TOC

Network Working Group	S. Bryant, Ed.
Internet-Draft	Cisco
Intended status: Standards Track	N. Sprecher, Ed.
Expires: June 26, 2010	Nokia Siemens Networks
	H. van Helvoort, Ed.
	Huawei
	A. Fulignoli, Ed.
	Ericsson
	Y. Weingarten
	Nokia Siemens Networks
	December 23, 2009

MPLS-TP Linear Protection
draft-weingarten-mpls-tp-linear-protection-05.txt

Abstract

The MPLS Transport Profile (MPLS-TP) being specified jointly by IETF and ITU-T includes requirements documents and framework documents. The framework documents define the basic architecture that is needed in order to support various aspects of the required behavior. This document addresses the functionality described in the MPLS-TP Survivability Framework document and defines a protocol that may be used to fulfill the function of the Protection State Coordination for linear protection, as described in that document.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunications Union Telecommunications Standardization Sector (ITU-T) effort to include an MPLS Transport Profile within the IETF MPLS and PWE3 architectures to support the capabilities and functionalities of a packet transport network as defined by the ITU-T.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction
    1.1. Contributing authors
2. Conventions used in this document
    2.1. Acronyms
    2.2. Definitions and Terminology
3. Protection switching logic
    3.1. Protection switching trigger mechanisms
    3.2. Protection switching control logical architecture
        3.2.1. PSC Status Module
4. Protection state coordination (PSC) protocol
    4.1. Transmission and acceptance of PSC control packets
    4.2. Protocol format

4.2.1. PSC Requests
4.2.2. Protection Type (PT)
4.2.3. Path fault identifier (FPath)
4.2.4. Active path indicator (Path)
4.3. Principles of Operation
4.3.1. PSC States
4.3.2. Failure or Degraded condition (Working path)
4.3.3. Lockout of Protection
4.3.4. Failure or Degraded condition (Recovery path)
4.3.5. Operator Controlled Switching
4.3.6. Recovery from switching
5. IANA Considerations
6. Security Considerations
7. Acknowledgements
8. References
8.1. Normative References
8.2. Informative References
§ Authors' Addresses

1. Introduction

TOC

As noted in the architecture for Multi-Protocol Label Switching Transport Profile (MPLS-TP) [TPFwk] (Bocci, M., Bryant, S., and L. Levrau, “A Framework for MPLS in Transport Networks,” July 2009.), the overall architecture framework for MPLS-TP is based on a profile of the MPLS and Pseudowire (PW) procedures as specified for the MPLS and (MS-)PW architectures defined in [RFC3031] (Rosen, E., Viswanathan, A., and R. Callon, “Multiprotocol Label Switching Architecture,” Jan 2001.), [RFC3985] (Bryant, S. and P. Pate, “Pseudowire Emulation Edge-to-Edge (PWE3) Architecture,” March 2005.) and [RFC5085] (Nadeau, T. and C. Pignataro, “Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires,” December 2007.). One of the basic survivability functions, pointed out by the Survivability Framework document [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” Feb 2009.), is that of simple and rapid protection switching mechanisms for Label Switched Paths (LSP) and Pseudo-wires (PW).

Protection switching is a fully allocated survivability mechanism. It is fully allocated in the sense that the route and bandwidth of the recovery path is reserved for a selected working path or set of working paths. It provides a fast and simple survivability mechanism, that allows the network operator to easily grasp the active state of the network, compared to other survivability mechanisms.

As specified in the Survivability Framework document [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” Feb 2009.), protection switching is applied to a protected domain. For the purposes of this document, we define the protected domain of a P2P LSP as consisting of two Label Switching Routers (LSR) and the transport paths that connect them. For a P2MP LSP the protection domain includes the root (or source) LSR, the destination (or sink) LSRs, and the transport paths that connect them.

In 1+1 unidirectional architecture as presented in [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” Feb 2009.), a recovery transport path is dedicated to each working transport path. Normal traffic is bridged (as defined in [RFC4427] (Mannie, E. and D. Papadimitriou, “Recovery Terminology for Generalized Multi-Protocol Label Switching,” Mar 2006.))and fed to both the working and the recovery transport entities by a permanent bridge at the source of the protection domain. The sink of the protection domain selects which of the working or recovery entities to receive the traffic from, based on a predetermined criteria, e.g. server defect indication. When used for bidirectional switching the 1+1 protection architecture must also support a Protection State Coordination (PSC) protocol. This protocol is used to help synchronize the decisions of both ends of the protection domain in selecting the proper traffic flow.

In the 1:1 architecture, a recovery transport path is dedicated to the working transport path of a single service. However, the normal traffic is transmitted only once, on either the working or the recovery path, by using a selector bridge at the source of the protection domain. A selector at the sink of the protection domain then selects the path that carries the normal traffic. Since the source and sink need to be coordinated to ensure that the selector bridge at both ends select the same path, this architecture must support a PSC protocol.

The 1:n protection architecture extends this last architecture by sharing the recovery path amongst n services. Again, the recovery path is fully allocated and disjoint from any of the n working transport paths that it is being used to protect. The normal data traffic for each service is transmitted only once, either on the normal working path for that service or, in cases that trigger protection switching (as defined in [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” Feb 2009.)), may be sent on the recovery path. It should be noted that in cases where multiple working path services have triggered protection switching that some services, dependent upon their Service Level Agreement (SLA), may not be transmitted as a result of limited resources on the recovery path. In this architecture there is a need for coordination of the protection switching, and in addition there is need for resource allocation negotiation. Due to the added complexity of this architecture, the procedures for this will be delayed to a different document and further study.

As was pointed out in the Survivability Framework [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” Feb 2009.) and highlighted above, there is a need for coordination between the end-points of the protection domain when employing bidirectional protection schemes. This is especially true when there is a need to maintain traffic over a co-routed bidirectional LSP. This document presents a protocol and a set of procedures for activating this coordination within the protection domain.

1.1. Contributing authors

TOC

Hao Long (Huawei)

2. Conventions used in this document

TOC

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] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.).

2.1. Acronyms

TOC

This draft uses the following acronyms:

DNR	Do not revert
FS	Forced Switch
GACH	Generic Associated Channel Header
LSR	Label Switching Router
MPLS-TP	Transport Profile for MPLS
MS	Manual Switch
P2P	Point-to-point
P2MP	Point-to-multipoint
PDU	Packet Data Unit
PSC	Protection State Coordination Protocol
PST	Path Segment Tunnel
SD	Signal Degrade

2.2. Definitions and Terminology

TOC

The terminology used in this document is based on the terminology defined in [RFC4427] (Mannie, E. and D. Papadimitriou, “Recovery Terminology for Generalized Multi-

SF	Signal Fail
SLA	Service Level Agreement
WTR	Wait-to-Restore

Protocol Label Switching," Mar 2006.) and further adapted for MPLS-TP in [SurvivFwk] (Sprecher, N., Farrel, A., and H. Shah, "Multi-protocol Label Switching Transport Profile Survivability Framework," Feb 2009.). In addition, we use the term LSR to refer to a MPLS-TP Network Element, whether it is a LSR, LER, T-PE, or S-PE.

3. Protection switching logic

TOC

3.1. Protection switching trigger mechanisms

TOC

The protection switching should be initiated in reaction to any of the following triggers:

- Server layer indication – if the MPLS-TP server layer detects a failure within its own layer, or due to a failure of its server layer (e.g. the physical layer) notifies the MPLS-TP layer that a failure has been detected.

- OAM signalling – if, for example, OAM continuity and connectivity verification tools detect that there is a loss of continuity or mis-connectivity or performance monitoring indicates a degradation of the utility of the working path for the current transport path. In cases of signal degradation, switching to the recovery path SHOULD only be activated if the recovery path can guarantee better conditions than the degraded working path.

- Control plane – if there is a control plane active in the network (either signaling or routing), it MAY trigger protection switching based on conditions detected by the control plane. If the control-plane is based on GMPLS [RFC3945] (Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," Oct 2004.) then the recovery process should comply with the process described in [RFC4872] (Lang, J., Papadimitriou, D., and Y. Rekhter, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery," May 2007.).

- Operator command – the network operator may issue commands that trigger protection switching. The commands that are supported include – Forced Switch, Manual Switch, Clear, Lockout of

Protection, (see definitions in [RFC4427] (Mannie, E. and D. Papadimitriou, “Recovery Terminology for Generalized Multi-Protocol Label Switching,” Mar 2006.)).

3.2. Protection switching control logical architecture

TOC

Protection switching processes the triggers described above together with the inputs received from the far-end LSR. These inputs cause the LSR to take certain actions, e.g. switching the Selector Bridge to select the working or recovery path, and to transmit different protocol messages.

```
+-------------+ Operator Command       Local PSC      +-----------+
|  External   |-----------------+   +-----------------| PSC Status|
|  Interface  |                 |   |   request   +---|   Module  |
+-------------+                 |   |             |   +-----------+
                                V   V             V Prot. Stat. ^
+----------+ Local OAM  +---------------+Highest +------------+ |
|   OAM    |----------->| Local Request |------->|  PSC Mess. | |
|  Module  |  request   |    logic      |local R.| Generator  | |
+----------+   +------->+---------------+        +------------+ |
+----------+   |                      |                 |       |
| Svr/CP   |---+         Highest local|request          |       |
+----------+                          V                 V       |
+-------------+             +-----------------+    PSC Message  |
| Remote Req. | Remote PSC  |  global Request |                 |
|  Receiver   |------------>|      logic      |                 |
+-------------+   Request   +-----------------+                 |
       ^                             |                          |
       |       Highest global request|                          |
       |                             V                          |
       |                    +-----------------+   PSC status    |
Remote PSC message          |    PSC Process  |-----------------+
                            |       logic     |--------> Action
                            |                 |
                            +-----------------+
```

Figure 1: Protection switching control logic

Figure 1 (Protection switching control logic) describes the logical architecture of the protection switching control. The Local Request logic unit accepts the triggers from the OAM, external operator commands, and from the local control plane (when present) and determines the highest priority request. This high-priority request is passed to both the PSC Message generator, that will generate the appropriate protocol message to be sent to the far-end LSR, and the Global Request logic, that will cross-check this local request with the information received from the far-LSR. The Global Request logic then processes these two PSC requests that determines the highest priority request that is passed to the PSC Process logic. The PSC Process logic uses this input to determine what actions need to be taken, e.g. switching the Selector Bridge, and the current status of the protection domain.

3.2.1. PSC Status Module

TOC

The PSC Control Logic must retain the status of the protection domain. The possible different states indicate the current status of the protection environment, and can be in one of three states:

- Normal (Idle) state – When both the recovery and the working paths are fully allocated and active, data traffic is being transmitted over the working path, and there are no trigger events reported within the domain.

- Protecting state – When either the working path has reported a signal failure (SF) or degradation of signal (SD), or the operator has issued an operator command and the data traffic has been redirected to the recovery path.

- Unavailable state – When the recovery path is unavailable, either as a result of reporting a SF or SD condition, or as a result of an administrative Lockout command.

This state may affect the actions taken by the control logic, and therefore, the PSC Status Module transfers the current status to the Local Request Logic.
See section 4.3.1 for details on what actions are affected by the PSC state.

TOC

4. Protection state coordination (PSC) protocol

Bidirectional protection switching, as well as unidirectional 1:1 protection, requires coordination between the two end-points in determining which of the two possible paths, the working or recovery path, is transmitting the data traffic in any given situation. When protection switching is triggered as described in section 3.1, the end-points must inform each other of the switch-over from one path to the other in a coordinated fashion.

There are different possibilities for the type of coordinating protocol. One possibility is a two-phased coordination in which the MEP that is initiating the protection switching sends a protocol message indicating the switch but the actual switch-over is performed only after receiving an 'Ack' from the far-end MEP. The other possibility is a single-phased coordination, in which the initiating MEP switches over to the alternate path and informs the far-end MEP of the switch, and the far-end MEP must complete the switch-over.

In the following sub-sections we describe the protocol messages that should be used between the two end-points of the protection domain.

For the sake of simplicity of the protocol, this protocol is based on the single-phase approach described above.

4.1. Transmission and acceptance of PSC control packets

TOC

The PSC control packets should be transmitted over the recovery path only. This allows the transmission of the messages without affecting the normal traffic in the most prevalent case, i.e. the idle state. In addition, limiting the transmission to a single path avoids possible conflicts and race conditions that could develop if the PSC messages were sent on both paths.

Any new PSC control packet must be transmitted immediately when a change in the transmitted status occurs.

When the PSC information is changed, three PSC packets should be transmitted as quickly as possible, so that fast protection switching would be possible. Transmission of three rapid packets allows for fast protection switching even if one or two PSC packets are lost or corrupted. The frequency of the first three packets and the separate frequency of the continual transmission is configurable by the operator. For protection switching within 50ms, the default interval of the first three PSC signals should be no larger than 3.3ms. PSC packets after the first three should be transmitted with an interval of 5 seconds.

If no valid PSC specific information is received, the last valid received information remains applicable. In the event a signal fail condition is detected on the recovery path, the received PSC specific information should be evaluated.

4.2. Protocol format

TOC

The protocol messages SHALL be sent over the GACH as described in [RFC5586] (Vigoureux,, M., Bocci, M., Swallow, G., Aggarwal, R., and D. Ward, “MPLS Generic Associated Channel,” May 2009.). There is a single channel type for the set of PSC messages, each message will be identified by the first field of the ACH payload as described below. PSC messages SHOULD support addressing by use of the method described in [RFC5586] (Vigoureux,, M., Bocci, M., Swallow, G., Aggarwal, R., and D. Ward, “MPLS Generic Associated Channel,” May 2009.). The following figure shows the format for the full PSC message.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|0 0 0 1|0 0 0 0|0 0 0 0 0 0 0 0|   MPLS-TP PSC Channel Code    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    ACH TLV Header                             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                    Addressing TLV                             +
:                              ...                              :
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
~                    PSC Control Packet                         ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: Format of PSC packet with a GACH header

Where:

- MPLS-TP PSC Channel Code is the GACH channel number assigned to the PSC = TBD

- The ACH TLV Header is described in [RFC5586] (Vigoureux,, M., Bocci, M., Swallow, G., Aggarwal, R., and D. Ward, “MPLS Generic Associated Channel,” May 2009.)

- The use of the Addressing TLV are for further study

*The following figure shows the format of the PSC Control message that is the payload for the PSC packet.

Editor's note: There is a suggestion that this format should be aligned with the format used by G.8031/G.8131/Y.1731 in ITU. The argument being that this would make it easier to pass review from ITU and allow easier transfer of technology.
The counter-argument is that the ITU format is based upon an attempt to find a common format for different functionality and therefore involves different fields that are not necessary for the protection switching. Defining a new dedicated format would make for a simpler and more intuitive protocol. End of editor's note.

```
0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|Ver|Request| PT|   FPath     |      Path        |    Reserved   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: Format of the PSC control packet

Where:

*Ver: is the version of the protocol, for this version the value SHOULD be 0.

*Request: this field indicates the specific PSC request that is being transmitted, the details are described in section 4.2.1

*PT: indicates the type of protection scheme currently supported, more details are given in section 4.2.2

*FPath: used to indicate the path that is reporting a failure condition, the possible values are described in section 4.2.3

*Path: used to indicate the currently active path, possible values are described in section 4.2.4

*Reserved: field is reserved for possible future use. These bits MUST be set to zero on transmission, and ignored upon reception.

TOC

4.2.1. PSC Requests

The Protection State Coordination (PSC) protocol SHALL support the following request types, in order of priority from highest to lowest:

*(1111) Clear

*(1110) Lockout protection

*(1101) Forced switch

*(0110) Signal fault

*(0101) Signal degrade

*(0100) Manual switch

*(0011) Wait to restore

*(0010) Do not revert (DNR)

*(0000) No request

See section 6.3 for a description of the operation of the different requests.

4.2.2. Protection Type (PT)

TOC

The PT field indicates the currently configured protection architecture type, this should be validated to be consistent for both ends of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values:

*11: 1+1 bidirectional switching

*10: 1:1 bidirectional switching

*01: 1+1 unidirectional switching

*00: 1:1 unidirectional switching

TOC

4.2.3. Path fault identifier (FPath)

The Fpath field of the PSC control SHALL be used only in a Signal fault (0101) or Signal degrade (0100) control packet. Its value indicates on which path the signal anomaly was detected. The following are the possible values:

- 0: indicates that the fault condition is on the Recovery path

- 1: indicates that the fault condition is on the Working path

- 2-255: for future extensions

4.2.4. Active path indicator (Path)

TOC

The Path field of the PSC control SHALL be used to indicate which path the source MEP is currently using for data transmission. The MEP should compare the value of this bit with the path that is locally selected for data transmission to verify that there is no inconsistency between the two end-points of the protected domain. If an inconsistency is detected then an alarm should be raised. The following are the possible values:

- 0: indicates that normal traffic is being transmitted on the Working path.

- 1: indicates the Recovery path is being used to transmit the normal traffic from the Working path.

- 2-255: for future extensions

4.3. Principles of Operation

TOC

In all of the following sub-sections, assume a protected domain between LSR-A and LSR-Z, using paths W (working) and R (recovery) as shown in figure 4.

```
+-----+ //=======================\\ +-----+
|LSR-A|//     Working Path        \\|LSR-Z|
|    /|                             |\    |
|  ?< |                             | >?  |
|    \|\\     Recovery Path       //|/    |
+-----+ \\=======================// +-----+

    |--------Protected Domain---------|
```

Figure 4: Protected domain

4.3.1. PSC States

TOC

4.3.1.1. Normal State

TOC

When the protected domain has no special condition in effect, the ingress LSR SHOULD forward the user data along the working path, and, in the case of 1+1 protection, the Permanent Bridge will bridge the data to the recovery path as well. The receiving LSR SHOULD read the data from the working path.

The ingress LSR MAY transmit a No Request PSC packet with the Path field set to 0 indicating that the normal data traffic should be read from the working path.

4.3.1.2. Protecting State

TOC

When the protection mechanism has been triggered and the protected domain has performed a protection switch, the domain is in the protecting state. In this state the normal data traffic is transmitted and received on the recovery path.

If the protection domain is currently in a protecting state, then the LSRs SHOULD NOT accept a Manual Switch request.

If the protection domain is currently in a protecting state, and a Forced Switch is requested then the normal traffic SHALL continue to be transmitted on the recovery path even if the original protection trigger is cleared, and the Forced Switch condition will be signalled by the PSC messages.

4.3.1.3. Unavailable State

TOC

When the recovery path is unavailable – either as a result of a Lockout operator command (see section 4.3.3), or as a result of a SF or SD detected on the recovery path (see section 4.3.4) – then the protection domain is in the unavailable state. In this state, the normal traffic is transmitted and received on the working path.
While in unavailable state any event that would trigger a protection switching SHOULD be ignored with the following exception – If a Signal Degrade request is received, then protection switching will be activated only if the recovery path can guarantee a better signal than the working path.
The protection domain will exit the unavailable state and revert to the normal state when, either the operator clears the Lockout command or the recovery path recovers from the signal fault or degraded situation. Both ends will resume sending the PCS packets over the recovery path, as a result of this recovery.

4.3.2. Failure or Degraded condition (Working path)

TOC

If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the working path that warrants invoking protection switching, then it SHOULD take the following actions:

- (For 1:1 protection) Switch all traffic for LSR-Z to the recovery path only.

- Transmit a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that normal traffic is now being transmitted on the recovery path.

- Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system.

When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions:

*Check priority of the request

*Switch all traffic addressed to LSR-A to the recovery path only (for 1:1 protection).

*Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being transmitted on the recovery path.

4.3.3. Lockout of Protection

TOC

If one of the LSRs (for example, LSR-A) receives a management command indicating that the protection is disabled, then it SHOULD indicate this to the far-end LSR (LSR-Z in this example) that it is not possible to use the recovery path. The following actions MUST be taken:

*Transmit a PCS control packet, using GACH, with the Request code set to Lockout of protection (1110), the Fpath set to 0, and the Path set to 0.

*All normal traffic packets should be transmitted on the working path only.

*Verify that the far-end LSR (for example LSR-Z) is forwarding the data packets on the working path. Raise alarm in case of mismatch.

*The PSC control logic should go into Unavailable state.

When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions:

*Check priority of request

*Switch all normal traffic addressed to LSR-A to the working path only.

*The PSC control logic should go into Unavailable status.

- *Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (Lockout of protection), the Fpath set to 0, and the Path set to 0, indicating that traffic is now being transmitted on the working path only.

4.3.4. Failure or Degraded condition (Recovery path)

TOC

If one of the LSRs (for example, LSR-A) detects a failure condition or a serious degradation condition on the recovery path, then it SHOULD take the following actions:

- *Begin transmission of a PCS control packet with the appropriate Request code (either Signal fault or Signal degrade), the Fpath set to 0, to indicate that the fault/degrade was detected on the recovery path, and the Path set to 0, indicating that traffic is now being forwarded on the working path. Note that this will actually reach the far-end if this is a unidirectional fault or recovery path is possibly in a degraded situation.
- *The PSC control logic should go into Unavailable state.
- *All traffic MUST be transmitted on the working path for the duration of the SF/SD condition.

When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has become Unavailable, it SHOULD perform the following actions:

- *Transmit all traffic on the working path for the duration of the SF/SD condition
- *The PSC Control logic should go into Unavailable state.

4.3.5. Operator Controlled Switching

TOC

If the management system indicated to one of the LSRs (for example LSR-A) that a switch is necessary, e.g. either a Forced Switch or a Manual Switch, then the LSR SHOULD switch the traffic to the recovery path and perform the following actions:

- *Switch all data traffic to the recovery path only.

- Transmit a PCS control packet, using GACH, with the appropriate Request code (either Manual switch or Forced switch), the Fpath set to 0, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being forwarded on the recovery path.

- Verify that LSR-Z replies with a PCS control packet indicating that it has switched to the recovery path. If this is not received after 2 PSC cycles then send an alarm to the management system.

When the far-end LSR (in this example LSR-Z) receives the PCS packet informing it that other LSR (LSR-A) has switched, it SHOULD perform the following actions:

- Check priority of the request

- Switch all normal traffic addressed to LSR-A to the recovery path only.

- Begin transmission of a PCS control packet, using GACH, with the appropriate Request code (either Manual switch of Forced switch), the Fpath set to 1, to indicate that the fault/degrade was detected on the working path, and the Path set to 1, indicating that traffic is now being forwarded on the recovery path.

4.3.5.1. Clearing operator commands

TOC

The operator may clear the switching condition by issuing a Clear request. This command will cause immediate recovery from the switch that was initiated by any of the previous operator commands, i.e. Forced Switch or Manual Switch. In addition, a Clear command after a Lockout Protection command should clear the Unavailable state and return the protection domain to the normal state.
If the Clear request is issued in the absence of a Manual Switch, Forced Switch, or Lockout protection, then it SHALL be ignored. In the presence of any of these commands, the Clear request SHALL clear the state affected by the operator command.

4.3.6. Recovery from switching

TOC

When the condition that triggered the protection switching clears, e.g. the cause of the failure condition has been corrected, or the operator clears a Manual Switch, then the protection domain SHOULD follow the following procedures:

- *If the network is configured for non-revertive behaviour, then the two LSRs SHOULD transmit DNR (Request code 0010) messages.
- *If the network is recovering from an operator switching command (in revertive mode), then both LSRs SHOULD return to using the working transport path and transmit No request (Request code 0000) messages.
- *If the network is recovering from a failure or degraded condition (in revertive mode), then the LSR that detects this recovery SHALL activate a local Wait-to-restore (WTR) timer (see section 4.3.6.1) to verify that there is not an intermittent failure. After the WTR expires, the LSR SHOULD return to using the working transport path and transmit No request (Request code 0000) messages.

4.3.6.1. Wait-to-restore timer

TOC

In revertive mode, in order to prevent frequent activation of protection switching due to an intermittent defect, the working transport path must become stable and fault-free before reverting to the normal condition. In order to verify that this is the case a fixed period of time must elapse before the normal traffic uses the working transport path. This period, called the Wait-to-restore (WTR) period, should be configurable by the operator in 1-minute intervals within the range 1-12 minutes. The default value is 5 minutes.
During this period, if a failure condition is detected on the working transport path, then the WTR timer is cleared and the normal traffic SHALL continue to be transported over the recovery transport path. If the WTR timer expires without being pre-empted by a failure, then the traffic SHOULD be returned to use the working transport path (as above).

5. IANA Considerations

TOC

To be added in future version.

TOC

6. Security Considerations

To be added in future version.

7. Acknowledgements

TOC

The authors would like to thank all members of the teams (the Joint Working Team, the MPLS Interoperability Design Team in IETF and the T-MPLS Ad Hoc Group in ITU-T) involved in the definition and specification of MPLS Transport Profile.

8. References

TOC

8.1. Normative References

TOC

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

8.2. Informative References

TOC

[RFC3031]	Rosen, E., Viswanathan, A., and R. Callon, “Multiprotocol Label Switching Architecture,” RFC 3031, Jan 2001.
[RFC3985]	Bryant, S. and P. Pate, “Pseudowire Emulation Edge-to-Edge (PWE3) Architecture,” RFC 3985, March 2005 (TXT).
[RFC5085]	Nadeau, T. and C. Pignataro, “Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires,” RFC 5085, December 2007 (TXT).
[TPFwk]	Bocci, M., Bryant, S., and L. Levrau, “A Framework for MPLS in Transport Networks,” ID draft-ietf-mpls-tp-framework-06.txt, July 2009.
[RFC5586]	Vigoureux,, M., Bocci, M., Swallow, G., Aggarwal, R., and D. Ward, “MPLS Generic Associated Channel,” RFC 5586, May 2009.
[RFC4427]	Mannie, E. and D. Papadimitriou, “Recovery Terminology for Generalized Multi-Protocol Label Switching,” RFC 4427, Mar 2006.
[SurvivFwk]	

	Sprecher, N., Farrel, A., and H. Shah, “Multi-protocol Label Switching Transport Profile Survivability Framework,” ID draft-ietf-mpls-tp-survive-fwk-02.txt, Feb 2009.
[RFC4872]	Lang, J., Papadimitriou, D., and Y. Rekhter, “RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery,” RFC 4872, May 2007.
[RFC3945]	Mannie, E., “Generalized Multi-Protocol Label Switching (GMPLS) Architecture,” RFC 3945, Oct 2004.

Authors' Addresses

TOC

Stewart Bryant (editor)
Cisco
United Kingdom
Email: stbryant@cisco.com

Nurit Sprecher (editor)
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon, 45241
Israel
Email: nurit.sprecher@nsn.com

Huub van Helvoort (editor)
Huawei
Kolkgriend 38, 1356 BC Almere
Netherlands
Phone: +31 36 5316076
Email: hhelvoort@huawei.com

Annamaria Fulignoli (editor)
Ericsson
Italy
Phone:
Email: annamaria.fulignoli@ericsson.com

Yaacov Weingarten
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon, 45241
Israel
Phone: +972-9-775 1827
Email: yaacov.weingarten@nsn.com