General comments:

1. The document contains too many repetitions and it is difficult to understand whether they are consistent to each other. For example protection switching triggers are described in section 1.1, 1.2 and 4.1.
2. The description does not appear to be fully consistent with the claimed scope. For example, general description (applicable to both LSP and PW) is not clearly decoupled from LSP specific description.
3. The introduction states that PW recovery is out of scope, but the draft includes section 7 Pseudowire Protection Considerations and in particular 7.2. Recovery in the Pseudowire Layer.
    a. Please clarify.
    b. If the Pseudowire is considered, what is it: a separate layer, a client signal, or a path?
    c. Also the titles of 7 and 7.2 are one example of the terms protection and recovery being used inconsistently.
4. The draft also describes fault isolation and fault reporting: However fault isolation is not required to initiate recovery action (and in some cases may delay recovery). The fault only needs to be isolated to the recovery domain that can (and will) take action. Fault isolation and reporting should be covered in the OAM framework.
5. The framework does not include any consideration of multiple faults or the impact of MTTR (mean time to repair) on availability.
6. The framework mentions various types of shared protection but does not make it clear what is being shared. The label, allocation of capacity on the server, reservation of server capacity.
   For example if two protection LSPs "share" a common server LSP but each has its own capacity reservation and hence both can be accommodated simultaneously, is this considered shared.
   If the reservation for capacity is shared between several protection paths we need a means to notify the other working paths if one of them is occupying the protection paths - i.e. they are in effect unprotected. The current text in 4.3.2 indicates some of the consequences but offers no guidance on how the action can be controlled.
7. The terminology in this document is not always consistent:
    a. The document is using both “connection” per G.805 and “transport path” per RFC 5654 terms. According to RFC 5654, the two terms are identical. It is simpler to read the document if only one term is used.
    b. The usage of the terms “defect”, “failure” and “degradation” is a little bit confusing. It is proposed to use the term “defect” to indicate a condition detected by an MPLS-TP node in case of “failure” or performance “degradations” events. As such, it is the detection of a “defect” that triggers protection switching in case of “failure” or performance “degradation”.
    c. The use of the terms “recovery”, “protection” and “restoration” is not always consistent (see also detailed comments).

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 08, 2010

N. Sprecher
Nokia Siemens Networks
A. Farrel
Old Dog Consulting
March 08, 2010

Multiprotocol Label Switching Transport Profile Survivability Framework

draft-ietf-mpls-tp-survive-fwk-04.txt

Abstract

Network survivability is the ability of a network to rapidly restore traffic delivery following ~~disruption or~~ failure or degradation of network resources. Survivability is critical to the delivery of guaranteed network services such as those subject to strict Service Level Agreements (SLAs) that place maximum bounds on the length of time the service may be degraded or unavailable.

The Transport Profile of Multiprotocol Label Switching (MPLS-TP) is a packet transport technology based on the MPLS data plane and re-using many aspects of the MPLS management and control planes.

This document provides a framework for the provision of survivability in an MPLS-TP network, describing recovery elements, types, methods and topological considerations. Survivability may be supported by control plane, management plane, and by Operations, Administration and Maintenance (OAM) functions to achieve data plane recovery. This document describes mechanisms for ~~protecting~~ recovering MPLS-TP Label Switched Paths (LSPs). Detailed consideration for the ~~protection~~ recovery of pseudowires in MPLS-TP networks is out of scope.

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication 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.

Comment [Italo Bus1]: Page: 1
The use of term “disruption” in addition to “failure or degradation” is not clear. Which cases are intended to cover under “disruption” that are not covered by “failure or degradation”?

Comment [Italo Bus2]: Page: 1
I understand this document covers both protection and restoration of LSPs.

Comment [Italo Bus3]: Page: 1
I understand that details for both protection and restoration of PWs are outside the scope of this document.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at
http://www.ietf.org/shadow.html.

This Internet-Draft will expire on May 13, 2010.

Table of Contents

Editors' Note:

This Informational Internet-Draft is aimed at achieving IETF Consensus before publication as an RFC and will be subject to an IETF Last Call.

[RFC Editor, please remove this note before publication as an RFC and insert the correct Streams Boilerplate to indicate that the published RFC has IETF Consensus.]

1. Introduction

Network survivability is the network's ability to rapidly ~~restore~~ recover traffic delivery following a failure or degradation of traffic delivery caused by a network fault or an attack on the network; it plays a critical role in the delivery of reliable services in transport networks. Guaranteed services in the form of Service Level Agreements (SLAs) require a resilient network that very rapidly detects facility or node degradation or failures, and immediately starts to ~~restore~~ recover network operations in accordance with the terms of the SLA.

Comment [Italo Bus4]: Page: 1
I think this is more consistent with the terms recovery, protection and restoration as defined in RFC 4427.

Comment [Italo Bus5]: Page: 1
I think this is more consistent with the terms recovery, protection and restoration as defined in RFC 4427.

The MPLS Transport Profile (MPLS-TP) is described in ~~[RFC5654] and~~ [MPLS-TP-FWK] and [MPLS-TP-p2mp-FWK]. MPLS-TP is designed to be consistent with existing transport network operations and management models, and to provide survivability mechanisms, such as protection and restoration. The function provided is intended to be similar to or better than that found in established transport networks which set a high benchmark for reliability. That is, it is intended to provide the operator with functions with which they are familiar through their experience with other transport networks, but this does not preclude additional techniques.

Comment [Italo Bus6]: Page: 1
RFC 5654 defines the requirements for MPLS-TP and not MPLS-TP itself. Moreover, the mpls-tp-framework draft does not describe the p2mp aspects of MPLS-TP.

This document provides a framework for MPLS-TP-based survivability that meets the recovery requirements specified in [RFC5664].

It uses the recovery terminology defined in [RFC4427] which draws heavily on [G.808.1]~~, and it refers to the requirements specified in [RFC5654]~~.

Comment [Italo Bus7]: Page: 1
The relationship between RFC 4427 and G.808.1 does not seem well defined. RFC 4427 claims to "borrow from the G.808.1", this paragraph says that RFC 4427 "draws heavily on G.808.1" while section 2 of this document claims that RFC 4427 is "consistent with G.808.1".

Comment [Italo Bus8]: I think that this framework should not only refer to the requirements but meet them

This document is a product of a joint Internet Engineering Task Force (IETF) / International Telecommunication Union Telecommunication 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.

1.1. Recovery Schemes

Various recovery schemes (for protection and restoration) and processes have been defined and analyzed in [RFC4427] and [RFC4428]. These schemes can also be applied in MPLS-TP networks to re-establish end-to-end traffic delivery within the agreed service level and to recover from "failed" or "degraded" transport entities. In the context of this document, transport entities are nodes, links, Label Switch Path (LSP)segments, concatenated LSP segments, and whole LSPs.

Recovery actions are ~~normally~~ initiated by the detection of a defect ~~or~~ (in case of failure or performance degradation), or by an external request (e.g., an operator request for manual control of protection switching).

[RFC4427] makes a distinction between protection switching and restoration mechanisms~~.~~ .

Protection switching makes use of pre-assigned capacity between nodes, where the simplest scheme has one dedicated protection entity for each working entity, while the most complex scheme has m protection entities shared between n working entities (m:n). ~~Protection switching may be either unidirectional or bidirectional; unidirectional meaning that each direction of a bidirectional connection is protection switched independently, while bidirectional means that both directions are switched at the same time even if the fault applies to only one direction of the connection.~~

Restoration uses any capacity available between nodes and usually involves re-routing. The resources used for restoration

may be pre-planned (i.e., predetermined, but not yet allocated to the recovery path) and recovery priority may be used as a differentiation mechanism to determine which services are recovered and which are not recovered or are sacrificed in order to achieve recovery of other services. ~~Restoration may also be either unidirectional or bidirectional.~~

Both protection and restoration mechanisms may be either unidirectional or bidirectional; unidirectional meaning that each direction of a bidirectional connection is recovered independently, while bidirectional means that both directions are recovered at the same time even if the fault applies to only one direction of the connection, also coordination of the recovery at both ends is required.

Both protection and restoration mechanisms may be either revertive or non-revertive as described in section 4.11 of [RFC 4427].

In general, protection actions are completed within time frames of tens of milliseconds, while automated restoration actions are normally completed in periods ranging from hundreds of milliseconds to a maximum of a few seconds, and protection is guaranteed while restoration is not.

1.2. Recovery Action Initiation

The recovery schemes described in [RFC4427] and evaluated in

Comment [Italo Bus9]: Page: 1
□ It seems easier to read if the description of unidirectional and bidirectional recovery is made generic and repetitions are avoided.

Comment [Italo Bus10]: Page: 1
□ Proposed to add the description of revertive/non-revertive operations.

Comment [Italo Bus11]: Page: 1
□ Proposed edits aimed at improving readability of the text.

Comment [Italo Bus12]: Page: 1
□ This title does not seem to be aligned with the content of this section. It is proposed to merge 1.1 and 1.2 and to restructure a bit the content (e.g., avoiding duplication and introducing general concepts first).

[RFC4428] are presented in the context of control plane-driven actions (such as the configuration of the protection entities and functions, etc.). The presence of a distributed control plane in an MPLS-TP network is optional, and the absence of such a control plane does not affect the ability to operate the network and to use MPLS-TP forwarding, Operations, Administration and Maintenance (OAM), and survivability capabilities. In particular, the concepts discussed in [RFC4427] and [RFC4428] refer to recovery actions in the data plane and are equally applicable in MPLS-TP with or without the use of a control plane.

Thus, some of the MPLS-TP recovery mechanisms do not depend on a control plane, and use MPLS-TP OAM mechanisms or management actions to trigger protection switching ~~across connections that were set up using management plane configuration~~actions. ~~These recovery mechanisms may be triggered by data plane events or by operator actions, and are based on MPLS-TP OAM fault management functions. "Fault management" in this context refers to failure detection, localization, and notification (where the term "failure" is used to represent both signal failure and signal degradation). The term "trigger" is used to indicate any event that may be used to cause an implementation to consider taking protection action.~~

The principles of MPLS-TP protection switching operation are similar to those described in [RFC4427] as the protection mechanism is based on the ability to detect certain defects in the transport entities within the recovery domain. The protection switching controller does not care which monitoring method is used, as long as it can be given information about the status of the transport entities within the recovery domain (e.g., OK, signal failure, signal degradation, etc.).

The protection switching operation is basically a data-plane capability and in the context of MPLS-TP it needs to be ensured that it is possible to switch over independently of the way the network is configured and managed~~.~~ .

All the MPLS and GMPLS protection mechanisms

are applicable in an MPLS-TP environment, and it should be possible also to provision and manage the related protection entities and functions defined in MPLS and GMPLS using a management plane.

In some protection switching schemes (such as bidirectional protection switching), it is necessary to coordinate the protection state between the edges of the recovery domain. An MPLS-TP ~~Protection State Coordination (PSC)~~Automatic Protection Switching (APS) protocol may be used as an in-band (i.e., data plane-based) control protocol to align both ends of the ~~protected~~ protection domain. ~~Control~~ When the MPLS-TP control plane is in use, control plane-based mechanism can also be used to coordinate the protection states between the edges of the protection domain.

1.3. Recovery Context

An MPLS-TP LSP may be subject to any or all of MPLS-TP link recovery, path segment recovery, or end-to-end recovery, where:

Comment [Italo Bus13]: Page: 1
Do all the concepts in RFC 4427 and RFC 4428 refer to recovery actions in the data plane?

Comment [Italo Bus14]: Page: 1
I think the statement is correct irrespectively on how the connections are setup.

Comment [GA15]: The term “trigger” should be used as a means to activate protection

Comment [Italo Bus16]: Page: 1
This text is a repetition of the recovery triggers described in section 1.1 and it is very hard to read and understand. It seems simpler to just remove it.

Comment [Italo Bus17]: Page: 1
The applicability of existing MPLS and GMPLS mechanisms is a complete different concept that the fact that protection switching is a data-plane capability.

Comment [Italo Bus18]: Page: 1
The statement is too vague. At least the references to the documents where these mechanisms are defined need to be added.

- MPLS-TP link recovery refers to the recovery of an individual link (and hence all or a subset of the LSPs routed over the link) between two MPLS-TP nodes.

- Segment recovery refers to the recovery of an LSP segment (i.e., segment and concatenated segment in the language of [RFC5654]) between two nodes.

- End-to-end recovery refers to the recovery of an entire LSP from its ingress to its egress node.

More than one of these recovery techniques may be configured concurrently by a single LSP for added resiliency.

Co-routed bidirectional MPLS-TP LSPs are defined such that both directions of the LSP follow the same route through the network. In this case the directions are often required by the operator to fate-share (that is, if one direction fails, both directions should cease to operate). This may also be the case for associated bidirectional LSPs where the two directions of the LSP take different paths through the network. This causes a direct interaction between the recovery processing affecting the two directions of an LSP such that both directions of the LSP are recovered at the same time (i.e., bidirectional recovery is a consequence of fate sharing).

The recovery scheme operating at the data plane level can function in a multi-domain environment (in the wider sense of a "domain" [RFC4726]); it can also protect against a failure of a boundary node in the case of inter-domain operation.

Comment [Italo Bus19]: Page: 1
What about server layer protection to recover an individual link? Is link recovery, recovering the link or a sub-set of the connections routed over the link?

Comment [Italo Bus20]: Page: 1
This text looks generic. Why is it describing only LSP cases?

Comment [Italo Bus21]: Page: 1
Unidirectional and bidirectional protection switching mechanisms are already introduced in section 1.1.

MPLS-TP recovery schemes are intended to protect client traffic as it is sent across the MPLS-TP network. ~~This document introduces protection and restoration techniques in general terms and then describes how they may be applied in the LSP layer to meet the requirements of the MPLS-TP recovery schemes [RFC5654]. Section 7 also provides an introduction to how the techniques may be applied in the pseudowire layer, but detailed consideration of pseudowires is out of scope. A description of the MPLS-TP LSP and pseudowire layers can be found in [MPLS-TP-FWK].~~

Comment [Italo Bus22]: Page: 1
This text repeats the content of section 1.4.

1.4. Scope of this Framework

This framework introduces the architecture of the MPLS-TP recovery domain and describes the recovery schemes in MPLS-TP (based on the recovery types defined in [RFC4427] as well as the principles of operation, recovery states, recovery triggers, and information exchanges between the different elements that sustain the reference model~~. The reference model is based on the MPLS-TP OAM reference model which is defined in [MPLS-TP-OAM]~~.

The framework also describes the qualitative levels of the survivability functions that can be provided, such as dedicated recovery, shared protection, restoration, etc. The level of recovery directly affects the service level provided to the end-user in the event of a network failure.

Comment [Italo Bus23]: The intent of the sentence is not clear. It seems easier to just remove it.

The general description of the functional architecture is applicable for both LSPs and pseudowires (PWs), however, PW recovery is only introduced in Section 7, and the details are out of scope for this document.

This framework applies to general ~~LSP~~ recovery schemes, but also to schemes that are optimized for specific topologies in order to handle protection switching in an efficient manner. ~~Recovery schemes for PWs are introduced in Section 7, but the details are for further study and will be addressed in a separate document.~~

This document takes into account the need for co-ordination of protection switching ~~at~~ across multiple layers and sub-layers (for readability we use the term "layer" to refer equally to layers and sub-layers). This allows an operator to prevent races and allows the protection switching mechanism of one layer to recover from the failure event ~~fix a problem~~ before invoking protection switching at another layer.

This framework also specifies the functions that must be supported by MPLS-TP to support the recovery mechanisms. MPLS-TP introduces a tool kit to enable recovery in MPLS-TP-based networks and to ensure that affected traffic is recovered in the event of a failure.

Generally, network operators aim to provide the fastest, most stable, and the best protection mechanism at a reasonable cost according to the requirements of the customers. The greater the level of protection, the greater the number of resources consumed ~~and so the higher the likely cost both to the operator and to the customer~~. It is therefore expected that network operators will offer a wide spectrum of service levels. MPLS-TP-based recovery offers the flexibility to select the recovery mechanism, choose the granularity at which traffic is protected, and also choose the specific types of traffic that are to be protected. With MPLS-TP-based recovery, it is possible to provide different levels of protection for different classes of service, based on their service requirements.

2. Terminology and References

The terminology used in this document is consistent with that defined in [RFC4427]. That RFC is, itself, consistent with [G.808.1].

However, certain protection concepts (such as ring protection) are not discussed in [RFC4427], and for those concepts, terminology in this document is drawn from [G.841].

Readers should refer to those documents for normative definitions. This document supplies brief summaries of some terms for clarity and to aid the reader, but does not re-define terms.

In particular, note the distinction and definitions made in [RFC4427] for the following three terms.

- Protection: re-establishing end-to-end traffic using pre-allocated resources.

Comment [Italo Bus24]: I understand these general recovery schemes are applicable to both LSPs and PWs (given statement above).

Comment [Italo Bus25]: Repeated text. It is proposed to remove the repetition.

Comment [Italo Bus26]: Cost considerations are better kept outside the scope of standard documents.

Comment [Italo Bus27]: Is “classes of service” here intended as QoS classes? If this is not the case, it would be better to rephrase “classes of service” with “connections”.

- Restoration: re-establishing end-to-end traffic using resources allocated at the time of need. Sometimes referred to as "repair" of a service, LSP, or the traffic.

- Recovery: a generic term covering both Protection and Restoration.

Note that the term "survivability" as used in [RFC5654] to cover the functional elements or "protection" and "restoration" which are collectively known as "recovery".

Important background information on survivability can be found in [RFC3386], [RFC3469] , [RFC4426], [RFC4427], and [RFC4428].

In this document, the following additional terminology is applied:

- Fault Management as defined in [MPLS-TP-NM-Framework].

- Defect ~~and failure are~~is used to indicate a condition detected in case of ~~both signal defects and~~ failures ~~, and signal~~ or degradation events.

- Trigger indicates any event that ~~may be used to cause an implementation to consider taking~~ initiates a recovery action.

- The acronym OAM is defined as Operations, Administration and Maintenance consistent with [OAM-SOUP].

- A Transport Entity is a node, link, ~~Label Switch Path (LSP)~~ segment, concatenated ~~LSP~~ segment, or whole connection (LSP or PW).

- A Working Entity is a transport entity that carries traffic during normal network operation.

- A Recovery Entity is a transport entity that is used to ~~restore~~recover and transport traffic when the working entity fails.

General terminology for MPLS-TP is found in [MPLS-TP-FWK] and [ROSETTA]. Background information on MPLS-TP requirements can be found in [RFC5654].

3. Requirements for Survivability

MPLS-TP requirements are presented in [RFC5654] and serve as a normative reference for the definition of all MPLS-TP function including survivability. Survivability is presented in [RFC5654] as playing a critical role in the delivery of reliable services, and the requirements for survivability are set out using the recovery terminology defined in [RFC4427].

These requirements are summarized below. Reference numbers refer to the requirements as presented in [RFC5654]. Readers should refer to [RFC5654] for the definitive list of requirements which is not replaced or superseded by the list provided here.

Comment [Italo Bus28]: The definition has been generalized to cover also the PW case.

3.1. General Requirements

- Protection and restoration mechanisms must be provided (56).
- Recovery techniques should be as similar as possible to those in existing transport networks (56A).
- Point-to-point (P2P) and point-to-multipoint (P2MP) recovery

techniques should be the same if possible (56B).

- Recovery must be applicable to links, ~~transport paths,~~ segments, concatenated segments, and end-to-end connections (LSPs and PWs) (57).
- Recovery objectives must be configurable to meet the SLA objectives of the services offered including rapid (sub-50ms) recovery, protection of all traffic on a path, and protection across multiple domains (58, 59).
- The recovery mechanisms should be applicable to any topology (60). See also Section 3.4 of this document.
- Recovery must be coordinated across network layers (61).
- Recovery and reversion must not "flap" (62).

Note that there is no requirement for support for extra traffic [RFC4427] ~~except in a ring where MPLS-TP must support the sharing of protection bandwidth in a ring by allowing best-effort traffic (108). This form of extra traffic may sometimes referred to as "non-preemptable unprotected traffic".~~

Comment [Italo Bus29]: Page: 1
Up to so far the document is using the term "connection" as per G.805 rather than "transport path" as per RFC 5654. It is important to have a consistent terminology at least across the whole draft.

Comment [Italo Bus30]: Page: 1
Extra traffic is not required in MPLS-TP in any topology. NUT and protection bandwidth sharing are different concepts than extra-traffic.

3.2. Requirements for Restoration

- The restored ~~and protected~~ paths must be able to share resources (70).
- Priorities must be available to control the order of restoration and to facilitate preemption during restoration (71, 72).
- Reversion must be supported (73).

3.3. Requirements for Protection

- MPLS-TP data plane protection must operate without regard to payload content (63).
- The following protection schemes must be supported:
  - reversion (64).
  - unidirectional and bidirectional 1+1 protection for P2P (65A, 65B).
  - unidirectional 1+1 protection for P2MP (65C).

* bidirectional 1:n protection for P2P (67A).

* unidirectional 1:n protection for P2MP (67B).

- It must be possible to share protection resources (66). This includes:
  * 1:n mesh recovery should be supported (68).
  * sharing of resources between protection paths that will not be required to protect the same fault (69).

3.4. Requirements for Survivability in Ring Topologies

- MPLS-TP recovery mechanisms may be optimized for specific topologies provided such optimizations interoperate with~~, and are as similar as possible to,~~ standard techniques to provide end-to-end recovery (91~~, 100~~) and be as similar as possible to those in general transport networks (100).
- Ring topologies support must include:
  * single ring (92).
  * interconnected rings (93).
  * connection of rings to arbitrary networks (99).
  * logical and physical rings (101).
- Traffic protection in rings must include:
  * unidirectional and bidirectional P2P paths (94).
  * unidirectional P2MP paths (95).
- Ring recovery techniques:
  * must default to bidirectional (102).
  * must support reversion as the default behavior (103).
  * must distinguish (to the operator and for the purpose of prioritized recovery actions) trigger mechanisms (104).
  * should protect against multiple failures (106B).
  * must support sharing of protection resources (109).
  * must prevent recovery flapping (107).

Comment [Italo Bus31]: Req-100 of RFC 5654 asks for ring protection similarity with other transport technology and not for ring protection similarity with linear protection.

- Ring protection mechanism scaling must include:
  - 1+1 and 1:1 protection switching 50 ms from the moment of fault detection in a network with a 16-node ring with less than 1200km of fiber (96).
  - independence from the number of LSPs crossing the ring (97).
  - good scaling behavior (performance, memory, etc.) with increases in the number of transport paths, the number of nodes on the ring, and the number of ring interconnects (98).
- It must be possible to disable protection mechanisms on selected links in a ring (105).
- MPLS-TP recovery mechanisms in a ring must support prioritization of recovery actions arising from different commands or triggers and for different protected entities (106A).

3.5. Triggers for Protection, Restoration, and Reversion

~~Recall that a "trigger" is defined as any event that may be used to cause an implementation to consider taking recovery action.~~

- Triggers must be supported from:
  - lower network layers (74).
  - MPLS-TP OAM (75).
  - the management plane (76).
  - the control plane (if present) (78).
- It must be possible to distinguish trigger sources and to prioritize recovery action requests (77, 79).

3.6. Management Plane Operation

- Support is required for preplanning, pre-calculation, and pre-provisioning of recovery paths and groups of paths (80, 81, 82, 85).
- External commands (controls) must allow the operator to activate, prevent, or test without activating, any recovery operation (83, 84).
- It must be possible to configure all aspects of recovery (86).

- It must be possible to monitor all aspects of recovery (87, 88).

3.7. Control Plane and In-band OAM

- If a control plane is used, it must be possible operate all aspects of recovery (89).

- In-band OAM must support administrative control and protection state coordination (90).

4. Functional Architecture

This section presents an overview of the elements of the functional architecture for survivability within an MPLS-TP network. The intention is to decompose the survivability components into separate items so that it can be seen how they may be combined to provide different levels of recovery to meet the requirements set out in the previous section.

4.1. Elements of Control

Recovery is achieved through specific actions taken to repair network resources or to redirect traffic onto paths that avoid failures in the network. Those actions may be triggered automatically by the MPLS-TP network nodes upon detection of a network defect ~~or failure~~, or may be ~~under direct the control of~~triggered by an operator. Automatic action may be enhanced by in-band (i.e., data-plane based) OAM mechanisms ~~for fault management and performance monitoring~~, or by in-band or out-of-band control plane signaling.

4.1.1. ~~Manual~~ Operator Control

The survivability behavior of the network as a whole, and the reaction of each ~~LSP~~ connection when a fault is reported, may be under operator control. That is, the operator may establish network-wide or local policies that determine what actions will be taken when different defects ~~or failures~~ are reported that affect different ~~LSPs~~connections. ~~At the same time, when a service request is made to cause the establishment of one or more LSPs in the network, the operator (or requesting application) may express a required or requested level of service, and this will be mapped to particular survivability actions taken before and during LSP setup, after the discovery of a defect or failure of network resources, and upon recovery of those resources.~~

It should be noted that it is unusual to present a user or customer with options directly related to recovery actions. Instead, the user/customer enters into an SLA with the network provider, and the network operator maps the terms of the SLA (for example for

guaranteed delivery, availability, or reliability) onto recovery schemes within the network.

The operator can also be given ~~manual~~ control of ~~survivability~~recovery actions ~~and events~~. For example, the operator may perform the following actions:

- enable or disable survivability function
- induce the simulation of a network fault
- force a switchover from a working path to a recovery path.

Comment [Italo Bus32]: This section is quite confusing. It is mixing operator's policies to configure recovery mechanisms on different connections (usually at connection setup) with the operator's commands used to control the recovery actions on already configured recovery mechanisms. It seems more appropriate to clearly separate these different aspects.

Comment [Italo Bus33]: Page: 1
Is this section applicable only to LSPs or in general to any type of connection (i.e., LSPs and PWs)?

Comment [Italo Bus34]: This text is quite difficult to read. It seems duplicated with the text just below. It is proposed to remove it.

Forced switchover may be done for network optimization purposes with minimal disturbance of services, such as when modifying protected or unprotected services, when replacing MPLS-TP network nodes, etc. In some circumstances, a fault may be reported to the operator and the operator may then select and initiate the appropriate recovery action.

Comment [Italo Bus35]: Page: 1
This text should be aligned with the whole set of operator's commands described in G.808.1.

4.1.2. Defect ~~and Failure~~-Triggered Actions

Survivability actions may be directly triggered by network defects ~~and failures~~. That is, the device that detects the defect ~~or failure~~ (for example, notification of an issue reported from equipment in a lower layer, a failure to receive an OAM Continuity message, or a reception of OAM message reporting ~~a defect or failure~~a network failure condition) may immediately perform a survivability action. ~~Recall that the terms "defect" and "failure" are used to represent both signal defect / failure and signal degradation.~~

This behavior can be subject to management plane or control plane control, but does not require any control~~,~~ or management ~~or data plane~~ message exchange to trigger the recovery action; the action is directly triggered by events in the data plane. Note, however, that coordination of recovery actions between the edges of the recovery domain may require message exchanges for some recovery functions or when performing a bidirectional recovery action.

4.1.3. OAM Signaling

OAM signaling refers to data plane OAM message exchanges ~~that are in band or closely coupled to the data channel. Such messages may be~~ used to detect and isolate faults or indicate a degradation in the operation of the network. ~~, but in~~ In this context we are concerned with the use of these messages to ~~control or~~ detect network defects that trigger survivability actions. This requires the instantiation of an MEG between the nodes at the edge of the recovery domain.

Data plane OAM ~~signaling~~ may also be used to coordinate recovery actions within

the protection domain.

4.1.4. Control Plane Signaling

Control plane signaling is responsible for setup, maintenance, and teardown of transport paths that are not under management plane control. The control plane may also be used to coordinate the detection and isolation, and reaction to network defects ~~and failures~~ pertaining to peer relationships (neighbor-to-neighbor, or end-to-end). Thus, control plane signaling may initiate and coordinate survivability actions.

Comment [Italo Bus36]: The role of control plane to initiate protection switching actions is not very clear. Is it assumed to be used to report from an intermediate node a server-layer fault or defect that affects the LSP/PW to be protected?

The control plane can also be used to distribute topology and resource-availability information. In this way, "graceful shutdown" [GR-SHUT] of resources may be effected by withdrawing them, and this

can be used as a stimulus to survivability action in a similar way to the reporting or discovery of a fault as described in the previous sections.

4.2. Elements of Recovery

This section describes the elements of recovery. These are the quantitative aspects of recovery; that is the pieces of the network for which recovery can be provided.

Note that the terminology in this section is consistent with [RFC4427]. Where the terms differ from those in [RFC5654] a mapping is provided.

4.2.1. Span Recovery

A span is a single hop between neighboring MPLS-TP nodes in the same network layer. A span is sometimes ~~incorrectly~~ referred to as a link, and this may cause some confusion between the concept of a data link and a traffic engineering (TE) link. LSPs traverse TE links between neighboring MPLS-TP nodes in the MPLS-TP network layer, however, a TE link may be provided by:

- a single data link
- a series of data links in a lower layer established as an LSP and presented to the upper layer as a single TE link
- a set of parallel data links in the same layer presented either as a bundle of TE links, or a collection of data links that, together, provide data link layer protection scheme.

Thus, span recovery may be provided by:

- selecting a different TE link from a bundle
- moving the TE link so that it is supported by a different data link between the same pair of neighbors
- re-routing the LSP in the lower layer.

Moving the protected LSP to another TE link between the same pair of neighbors is a form of segment recovery and is described in Section 4.2.2.

4.2.2. Segment Recovery

Comment [Italo Bus37]: Is this description applicable in general or only to LSPs?

An LSP segment is one or more continuous hops on the path of the LSP. [RFC5654] defines two terms. A "segment" is a single hop on the path of an LSP, and a "concatenated segment" is more than one hop on the path of an LSP. In the context of this document, a segment covers both of these concepts.

~~A PW segment refers to a Single Segment PW (SS-PW) or to a single segment of a multi-segment PW (MS-PW) that is set up between two PE~~

~~devices (i.e., T-PE and S-PE, S-PE and S-PE, or S-PE and T-PE). As indicated in Section 1, the recovery of PWs and PW segments is out of scope of this document, but see Section 7.~~

~~LSP s~~Segment recovery involves redirecting or copying of traffic at the source end of a segment ~~of an LSP~~ onto an alternate path to the other end of the segment. According to the required level of recovery (described in Section 4.3), this redirection may be onto a pre-established ~~LSP~~ segment, through re-routing of the protected segment, or by tunneling the protected ~~LSP~~ segment through a "bypass" LSP. For details on recovery mechanisms, see Section 4.4.

Note that protecting ~~an LSP~~a connection against the failure of a node within the recovery domain requires the use of segment recovery, while a link failure could be protected using span or segment recovery.

4.2.3. End-to-End Recovery

End-to-end recovery is a special case of segment recovery where the protected ~~LSP~~ segment is the whole of the ~~LSP~~connection. End-to-end recovery may be provided as link-diverse or node-diverse recovery where the recovery path shares no links or no nodes with the working path. Note that node-diverse paths are necessarily link-diverse, and that full, end-to-end node-diversity is ~~required~~ necessary, although not sufficient as per section 4.7.2, to ~~guarantee recovery~~avoid a single point of failure.

Comment [Italo Bus38]: Is this description applicable in general or only to LSPs?

Comment [Italo Bus39]: When there is no full e2e diversity it is still possible to guarantee recovery in some cases (e.g., when the failures do not affect the so called “single points of failure”). Moreover, as described in section 4.7.2, full e2e node diversity is not a sufficient condition for avoiding single point of failures.

4.3. Levels of Recovery

This section describes the qualitative levels of survivability function that can be provided. The level of recovery offered has a direct effect on the service level provided to the end-user in the event of a network ~~fault~~failure or performance degradation. This will be observed as the amount of data lost when a network fault occurs, and the length of time to recover connectivity.

In general there is a correlation between the ~~service~~ recovery level (i.e., the rapidity of recovery and reduction of data loss) and the ~~cost to~~amount of resources used within the network; better service levels require pre-allocation of resources to the recovery paths, and those resources cannot be used for other purposes if high quality recovery is required. ~~Thus, "cost" in this case may be measured as the financial cost of providing resources for the recovery scheme, or the financial loss from dedicating resources to the recovery scheme such that they cannot be used to draw new revenue.~~

Comment [Italo Bus40]: Cost considerations are better kept outside the scope of standard documents.

Comment [Italo Bus41]: Cost considerations are better kept outside the scope of standard documents.

Sections 6 and 7 of [RFC4427] provide a full breakdown of protection and recovery schemes. This section summarizes the qualitative levels available.

4.3.1. Dedicated Protection

In dedicated protection, the resources for the recovery entity are pre-assigned for use only by the protected ~~service~~connection. This will clearly be the case in 1+1 protection, and ~~may~~ is also be the case in 1:1 protection ~~where~~ because extra traffic ~~(see Section 4.3.3)~~ is not supported.

Note that in the use of protection tunnels (see Section 4.4.3) resources may also be dedicated to protecting a specific ~~service~~LSP. In some cases (one-for-one protection) the whole of the bypass tunnel may be dedicated to provide recovery for a specific LSP, but in other cases (such as facility backup) a subset of the resources of the bypass tunnel may be pre-assigned for use to recover a specific ~~service~~LSP. However, as described in Section 4.4.3, the bypass tunnel approach can also be used for shared protection (Section 4.3.2)~~, to carry extra traffic (Section 4.3.3),~~ or, without reserving resources, to achieve best-effort recovery.

4.3.2. Shared Protection

In shared protection, the resources for the recovery entities of several services are shared. These may be shared as 1:n or m:n, and are shared on individual links. Link-by-link resource sharing may be managed and operated on LSP segments, on PW segments, or on end-to-

end transport path (LSP or PW). Note that there is no requirement for m:n recovery in the list of MPLS-TP requirements documented in [RFC5654].

Where a bypass tunnel is used (Section 4.4.3), the tunnel might not have sufficient resources to simultaneously protect all of the paths to which it offers protection, so that if they were all affected by network defects and failures at the same time, they would not all be recovered. Policy would dictate how this situation is handled: it might be that some individual paths would be protected while others would simply fail; it might be that the traffic for some paths would be guaranteed, but other traffic would be treated as best effort with the risk of packets being dropped~~; or it might be that protection would not be attempted~~.

Comment [Italo Bus42]: Isn't this the case where some path are protection while others would simply fail?

Shared protection is a trade-off between the dedication of ~~expensive~~ network resources to protection that is not required most of the time~~, and~~; the risk of unrecoverable services in the event of multiple network defects or failures;~~. Rapid~~ and rapid recovery that can be achieved with dedicated protection, but is delayed by message exchanges in the management, control, or data planes for shared protection. This means that there is also a trade-off between rapid recovery and the reduction of network ~~cost achieved by sharing protection~~ resources. Shared protection might not meet the protection speed requirements in some cases, but may still be faster than restoration.

These trade-offs may be somewhat mitigated by:

- adjusting the value of n in 1:n protection
- using m:n protection for some value of m > 1

- by establishing new protection paths as each available protection path is put into use.

4.3.3. Extra Traffic

~~Network resources allocated for protection represent idle capacity during the time that recovery is not actually required, and can be utilized by carrying other traffic referred to as "extra traffic".~~

~~Note that extra traffic does not need to start or be terminated at the ends of the entity (e.g. LSP) that it uses.~~

~~When a network resource that is carrying extra traffic is required for recovery of the protected traffic from the failed working path, the extra traffic is disrupted -- essentially it is pre-empted by the recovery LSP. This may require additional message exchanges in the~~

~~management, control, or data planes, and that may mean that recovery could be delayed. Thus the benefits of carrying extra traffic must be weighed against the disadvantage of delayed recovery, additional network overhead, and the impact to the services the extra traffic supports.~~

~~Note that extra traffic is not protected by definition, but may be restored.~~

~~Extra traffic is not supported on dedicated protection resources used for 1+1 protection (Section 4.3.1) by definition, but can be supported in other protection schemes including shared protection (Section 4.3.2) and tunnel protection (Section 4.4.3).~~

In MPLS-TP support for extra traffic is not required.

Best effort traffic should not be confused with extra traffic. Best effort traffic is such that the network does not provide any guarantees of data delivery, and the user is not given any guarantee of quality of service (e.g., in terms of jitter, packet loss, delay, etc.). Best effort traffic depends on the current traffic load, but extra traffic can have quality guarantees up until it is preempted by the need to use resources for recovery. At such a time, the extra traffic may be completely displaced, may be treated as best effort, or itself be recovered (for example, by restoration techniques).

~~Note that in MPLS-TP support for extra traffic is not required except in ring topologies (Section 3 and [RFC5654]).~~

4.3.4. Restoration

This section refers to LSP restoration. Restoration for PWs is out of scope of this document (but see Section 7).

Restoration represents the most effective use of network resources as no resources are reserved for recovery. However, restoration requires computation of a new path and activation of a new LSP (through the management or control plane). These steps can take much more time than is required for recovery using protection techniques.

Comment [HvH43]: This statement contradicts the statement in the last paragraph (starting with Note that)

Furthermore, there is no guarantee that restoration will be able to recover the service. It may be that all suitable network resources are already in use for other LSPs so that no new path can be found. This problem can be partially mitigated by the use of LSP setup priorities so that recovery LSPs can pre-empt existing LSPs of low priority.

Additionally, when a network defect or failure occurs, multiple LSPs may be disrupted by the same event. These LSPs may have been established by different Network Management Stations (NMSes) or

signaled by different head-end MPLS-TP nodes, and this means that multiple points in the network will be trying to compute and establish recovery LSPs at the same time. This can lead to contention for resources within the network, causing recovery failures and meaning that some recovery actions must be retried resulting in even slower recovery times for some services.

Both hard and soft LSP restoration may be supported. In hard LSP restoration, the resources of the working LSP are released before the full establishment of the recovery LSP (i.e., break-before-make). In soft LSP restoration, the resources of the working LSP are released after the full establishment of an alternate LSP (i.e., make-before-break). Note that, in the case of reversion (Section 4.3.5) the resource of the working LSP are not released.

Note that the restoration resources may be pre-calculated and even pre-signaled before the restoration action starts, but not pre-allocated. This is known as pre-planned LSP restoration. The complete establishment/activation of the restoration LSP occurs only when the restoration action starts. The pre-planning may happen periodically to have the most accurate information about the available resources in the network.

Comment [HvH44]: if the restoration path is pre signalled how is this different from protection

4.3.5. Reversion

After a service has been recovered so that traffic is flowing on the recovery LSP, the defective network resource may be replaced. The traffic can be redirected back on to the original working LSP (called "reversion"), or to left it where it is on the recovery LSP ("non-revertive" behavior).

It should be possible to specify the reversion behavior of each service, and this might even be configured for each recovery instance.

In the non-revertive mode an additional operational option exists where ~~protection~~ recovery roles are switched so that the recovery LSP becomes the working LSP, and the previous working path (or the resources used by the previous working path) are used for recovery in the event of a further fault.

In revertive mode it is important to prevent excessive swapping between working and recovery paths in the case of an intermittent defect. This can be addressed by the use of a reversion delay timer (the Wait To Restore timer) that controls the length of time to wait

following the repair of the fault on the original working path before performing reversion. It should be possible for an operator to configure this timer per ~~LSP~~connection, and a default value should be defined.

4.4. Mechanisms for Protection

The purpose of this section is to describe in general (MPLS-TP non-specific) terms the mechanisms that can be used to provide protection. As indicated above, while the functional architecture applies to both LSPs and PWs, the mechanism for recovery described in this document refers to LSPs and LSP segments only. Recovery mechanisms for pseudowires and pseudowire segment are for further study and will be described in a separate document (see also Section 7).

4.4.1. Link-Level Protection

Comment [HvH45]: This section is confusing, if the link is protected how can only some subset of the LSPs being carried be protected? This appears to be LSP protection not link protection.

Link-level protection refers to two paradigms: (1) where the protection is provided in a lower network layer, and (2) the protection is provided by the MPLS-TP link layer.

Note that link-level protection mechanisms do not protect the nodes at each end of the entity (e.g., a link or span) that is protected. End-to-end or segment protection should be used in conjunction to link-level protection to protect against a failure of the edge nodes.

Link-level protection offers the following levels of protections:

- Full protection, where a dedicated protection entity (e.g., a link or span) is pre-established to protect a working entity. When the working entity fails, the protected traffic is switched onto the protecting entity. In this scenario, all LSPs carried over the working entity are recovered (in one protection operation) when there is a failure condition. This is referred to in [RFC4427] as "bulk recovery".

- Partial protection, where only a subset of the LSPs or traffic carried over a given entity is recovered when there is a failure condition. The decision as to which LSPs will be recovered and which will not, depends on local policy.

~~When there is no failure on the working entity, the protection entity may transport extra traffic which may be preempted when protection switching occurs.~~

~~As with recovery in layered networks, a protection mechanism at the lower layer needs to be coordinated with protection actions at the upper layer in order to avoid race conditions. In general, this is arranged to allow protection actions to be performed in the lower layer before any attempt is made to perform protection actions in the upper layer.~~

Comment [Italo Bus46]: This text is in the scope of section 4.7.

A protection mechanism may be provided at the MPLS-TP link layer (which connects two MPLS-TP nodes). Such a mechanism can make use of the procedures defined in [RFC5586] to set up in-band communication channels at the MPLS-TP ~~link~~ section level and use these channels to monitor the health of the MPLS-TP link and coordinate the protection states between the ends of the MPLS-TP link.

Comment [Italo Bus47]: This text looks like a requirement. Which mechanism is provided?

4.4.2. Alternate Paths and Segments

The use of alternate paths and segments refers to the paradigm whereby protection is performed in the same network layer as the protected LSP either for the entire end-to-end LSP or for a segment of the LSP. In this case, hierarchical LSPs are not used - compare with Section 4.4.3.

Different levels of protection may be provided:

- Dedicated protection, where a dedicated entity (e.g., LSP or LSP segment) is fully pre-established to protect a working entity (e.g., LSP or LSP segment). When there is a failure condition on the working entity, the traffic is switched onto the protection entity. Dedicated protection may be performed using 1:1 or 1+1 linear protection schemes. When the failure condition is eliminated, the traffic may revert to the working entity. This is subject to local configuration.

- Shared protection, where one or more protection entity is pre-established to protect against a failure of one or more working entities (1:n or m:n).

When the fault condition on the working entity is eliminated, the traffic should revert back to the working entity in order to allow other related working entities to be protected by the shared protection resource.

4.4.3. Protection Tunnels

A protection tunnel is a hierarchical LSP that is pre-provisioned in order to protect against a failure condition along a sequence of spans in the network. We call such a sequence, a network segment. A failure of a network segment may affect one or more LSPs that transit the network segment.

When there is a failure condition in the network segment (detected either by OAM on the network segment, or by OAM on a concatenated segment of one of the LSPs transiting the network segment), one or more of the protected LSPs are switched over at the ingress point of the network segment and transmitted over the protection tunnel. The

way to realize this uses label stacking. Label mapping may be an option as well.

Different levels of protection may be provided:

- Dedicated protection, where the protection tunnel has resource reservations sufficient to provide protection for all protected LSPs without service degradation.

- Partial protection, where the protection tunnel has resources to protect some of the protected LSPs, but not all of them simultaneously. Policy would dictate how this situation is handled: it might be that some individual LSPs would be protected while others would simply fail; it might be that the traffic for some LSPs would be guaranteed, but traffic for other LSPs would be treated as best effort with the risk of packets being dropped; or it might be that protection would not be attempted.

4.5. Recovery Domains

Comment [Italo Bus48]: This section provides a good description of the recovery architecture. It would be better to read this text at the beginning of section 4.

Protection and restoration are performed in the context of a recovery domain. A recovery domain is defined between two or more recovery reference endpoints which are located at the edges of the recovery domain and bounds the element on which recovery can be provided (as described in Section 4.2 above). This element can be end-to-end path, a segment, or a span.

The case of an end-to-end path can be observed as a special case of a segment, and the ingress and the egress LERs serve as the recovery reference end-points.

In this simple case of a P2P protected entity, exactly two endpoints reside at the boundary of the Protection Domain. An LSP can enter through exactly one reference endpoint and exit the recovery domain through another reference endpoint.

In the case of unidirectional P2MP, three or more endpoints reside at the boundary of the Protection Domain. One of the endpoints is referred to as source/root and the other ones are referred to as sinks/leaves. An LSP can enter the recovery domain through the root point and exit the recovery domain through the leaves points.

The recovery mechanism should restore interrupted traffic due to a facility (link or node) fault within the recovery domain. Note that a single link may be part of several recovery domains. If two recovery domains have any links in common, then one recovery domain must be contained with the other. This can be referred to as nested recovery domains. The boundaries of recovery domains may coincide, but

recovery domains must not ~~intersect~~overlap.

Note that the edges of a recovery domain are not protected and unless the whole domain is contained within another recovery domain, the edges form a single point of failure.

A recovery group is defined within a recovery domain and it consists of a working (primary) entity and one or more recovery (backup) entities which reside between the endpoints of the recovery Domain. In order to guarantee protection in all situations, a necessary but not sufficient condition (see section 4.7.2) is to pre-provision a dedicated recovery entity ~~should be pre-provisioned~~ using disjoint resources in

the recovery domain in order to protect against a failure of a working entity.

The method used to monitor the health of the recovery element is outside the scope of this document. The endpoints that are responsible for the recovery action must receive the information on its condition. The condition of the recovery element may be 'OK', 'failed', or 'degraded'.

When the recovery operation is to be triggered by ~~an~~ OAM ~~FM or PM~~mechanisms ~~indication~~, an OAM Maintenance Entity Group must be defined for each of the working and protection entities.

The recovery entities and functions in a recovery domain can be configured using a management plane or a control plane. A management plane may be used to configure the recovery domain by setting the reference points, the working and recovery entities, and the recovery type (e.g., 1:1 bidirectional linear protection, ring protection, etc.). Additional parameters associated with the recovery process may also be configured. For more details, see Section 6.1.

When a control plane is used, the ingress LERs may communicate with the recovery reference points requesting that protection or restoration be configured across a recovery domain. For details, see Section 6.5.

Cases of multiple interconnections between distinct recovery domains actually just create a hierarchical arrangement of recovery domains as a single top-level recovery domain is created from the concatenation of two recovery domains that have multiple interconnections. In this case, recovery actions may be taken both in the individual lower-level recovery domains to protect any LSP segment that crosses the domain, and within the higher-level recovery domain to protect the longer LSP segment that traverses the higher-level domain.

Comment [Italo Bus49]: This issue has been already discussed in section 4.2.3. However, this looks like a better location as this consideration does not apply only to e2e recovery.

4.6. Protection in Different Topologies

As described in the requirements listed in Section 3 and detailed in [RFC5654], the recovery techniques used may be optimized for different network topologies if the performance of those optimized mechanisms is significantly better than the performance of the generic ones in the same topology.

It is required ([RFC5654] R91) that such mechanisms interoperate with the mechanisms defined for arbitrary topologies to allow end-to-end protection and to allow consistent protection techniques to be used across the whole network.

This section describes two different topologies and explains how recovery may be markedly different in those different scenarios. ~~It also introduces the concept of a recovery domain and shows how end-to-end survivability may be achieved through a concatenation of recovery domains each providing some level of recovery in part of the network.~~

Comment [Italo Bus50]: This is already well described in section 4.5.

4.6.1. Mesh Networks

Linear protection is a protection switching mechanism ~~of protection and protection state coordination~~ that provides a fast and simple ~~protection switching~~recovery ~~that fits best in mesh networks~~ in any network topology since it can operate between any pair of points within the network. It can protect against ~~a defect or failure~~ a failure or performance degradation that may happen on a~~n~~ node, a span, an LSP segment, or an end-to-end LSP. Linear protection provides a clear indication of the protection status.

Linear protection operates in the context of a Protection Domain. A Protection Domain is a special case of a Recovery Domain ~~[RFC4427]~~ (see section 4.5) that applies to the linear protection function~~.~~.

A Protection Domain is composed of the following architectural elements:

- A set of end points which reside at the boundary of the Protection Domain. In this simple case of 1:n or 1+1 P2P protection, exactly two endpoints reside at the boundary of the Protection Domain. In each transmission direction one of the end points is referred to as a source and the other one is referred to as a sink. In the case of unidirectional P2MP protection, three or more endpoints reside at the boundary of the Protection Domain. One of the endpoints is referred to as source/root and the other ones are referred to as sinks/leaves.

- A Protection Group which consists of ~~a~~ one or more working (primary) path~~s~~ and one ~~or more~~ protection (backup) path~~s~~ which run between the endpoints of the Protection Domain. ~~In order to guarantee protection in all situations, a dedicated protection path should be pre-provisioned to protect against a defect or failure of a working path (i.e., 1:1 or 1+1 protection schemes). Also the working and the protection paths should be disjoint, i.e.,, the physical routes of the working and the protection paths should have complete physical diversity.~~

Note that if the resources of the protection path are less than those of the working path, the protection path may not have sufficient resources to protect the traffic of the working path.

As mentioned in Section 4.3.2, the resources of the protection path may be shared as 1:n. In such a case, the protection path ~~might~~ cannot ~~have sufficient resources to~~ simultaneously protect all of the working paths ~~that may be affected by fault conditions~~ at the same time.

For P2P bidirectional paths, both unidirectional and bidirectional protection switching is supported. In bidirectional protection switching, in

Comment [Italo Bus51]: Linear protection mechanisms are applicable to any topology. Only shared mesh protection and ring protection seems to be the optimized mechanisms for specific topologies (i.e., mesh and ring topologies respectively).

Comment [Italo Bus52]: With 1:1 and 1+1 we have one working and one protection path; with 1:n we have n working paths and 1 protection paths. Only with m:n (not currently required for MPLS-TP) we have n working paths and m protection paths.

Comment [Italo Bus53]: This concept has been already explained before and it is quite generic and not specific to linear protection. It is better to remove the duplication.

Comment [Italo Bus54]: What is the purpose of setting up a protection path that does not have enough resources to protect its working path?

Comment [Italo Bus55]: This is true also when multiple working paths need to be protected at the same because of performance degradation or because of operator's commands.

Comment [Italo Bus56]: I do not think that bidirectional protection is applicable to unidirectional p2p paths.

~~the event of~~when a defect ~~or failure,~~ is detected on one direction, the protection actions are taken in both directions ~~(even when the fault is unidirectional)~~. This requires some level of coordination of the protection state between the endpoints of the protection domain.

In unidirectional protection switching, the protection actions are taken only in the affected direction.

Revertive and non-revertive operations (see section 4.3.5) are provided as network operator options.

4.6.1.1. Protection Schemes in Mesh Topologies

Linear protection supports the protection schemes described in the following sub-sections.

4.6.1.1.1. 1:n Linear Protection

In the 1:1 scheme, a protection path is allocated to protect against a defect or failure or degradation in a working path. As described above, in order to guarantee protection, the protection entity should support the full capacity and bandwidth~~, but it may be configured (for example, because of limited availability of network resources) to offer a degraded service compared to the working entity~~.

Figure 1 presents 1:1 protection architecture. In normal conditions the data traffic is transmitted over the working entity and the protection entity is in an idle state (OAM may be running on the protection entity to verify its state). Normal conditions are

defined when there is no defect, failure, or degradation on the working entity and there is no administrative configuration or request that cause traffic to transmit over the protection entity.

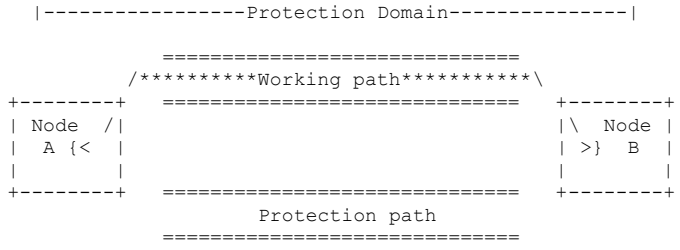

Figure 1: 1:1 Protection Architecture

Upon a ~~fault~~ defect condition (~~defect,~~ failure, or degradation) along the working entity or a specific administrative request, the traffic is switched over to the protection entity.

Note that in the non-revertive behavior (see Section 4.3.5), ~~the old working entity becomes the protection entity and extra traffic can be transmitted over the protection entity. This can happen~~ after the

Comment [GA57]: Odd name for a section that describes linear protection, should it be Linear Protection Schemes?

Comment [Italo Bus58]: It is not clear what is the value to support a complex scheme like linear protection to offer a degraded services when protection switching occurs.

conditions causing the switchover has/have been cleared, the traffic continues to flow on the protection path but the working and protection roles are not switched.

In each transmission direction, the source of the protection domain bridges the traffic into the appropriate entity and the sink selects the traffic from the appropriate entity. The source and the sink need to coordinate the protection states to ensure that the bridging and the selection are done to and from the same entity. For that sake a signaling coordination protocol (either data-plane in-band signaling protocol or a control-plane based signaling protocol) is needed.

In bidirectional protection switching, both ends of the protection domain switch to the protection entity (even when the fault is unidirectional). ~~This requires a protocol to try and coordinate the protection state between the two end points of the Protection Domain.~~

When there is no defect or failure, the bandwidth resources of the idle entity may be used ~~for less priority traffic~~by best-effort traffic. When protection switching is performed, ~~the lower priority traffic may be pre-empted by the protected traffic by tearing down the lower priority LSP, by reporting a fault on the lower priority LSP, or by treating the lower priority traffic as~~ best effort ~~and discarding it when~~ will be discarded if there is congestion.

Comment [Italo Bus59]: In linear protection it is actually the opposite case: the traffic continues to flow on the protection path but the working and protection roles are not switched. This is applicable also to the 1+1, so this note is more appropriate if moved just before the beginning of section 4.6.1.1.

Comment [Italo Bus60]: As written above, this is needed because of the 1:1 nature of the mechanism.

Comment [Italo Bus61]: Extra-traffic is not required/supported.

In the general case of 1:n linear protection, one protection entity is allocated to protect n working entities. ~~The protection entity might not have sufficient resources to simultaneously protect all of the working entities that may be affected by fault conditions at the same time~~In this case, in order to guaranteed protection, the protection entity should support enough capacity and bandwidth to protect any of the n working entities.

In case of defects or failures along multiple working entities, priority should be set as to which entity is protected. The protection states between the edges of the Protection Domain should be fully coordinated to ensure consistent behavior. As explained above in section revertive behavior it is recommended to use reversion when 1:n is supported.

Comment [GA62]: Section number is missing

4.6.1.1.2. 1+1 Linear Protection

In the 1+1 protection scheme, a fully dedicated protection entity is allocated.

As depicted in figure 2, data traffic is copied and fed at the source to both the working and the protection entities. The traffic on the working and the protection entities is transmitted simultaneously to the sink of the Protection Domain, where the selection between the working and protection entities is made (based on some predetermined criteria).

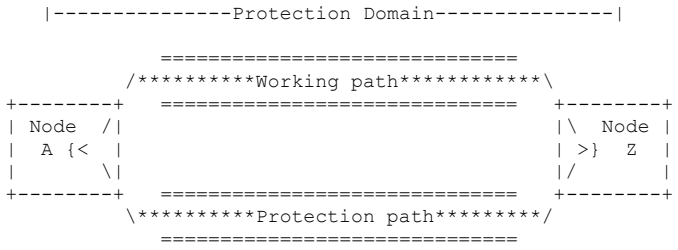

Figure 2: 1+1 Protection Architecture

Note that control traffic between the edges of the Protection Domain (such as OAM or control protocol to coordinate the protection state, etc.) may be transmitted on a different entity than the one used for the protected traffic. These packets should not be discarded by the sink.

In 1+1 unidirectional protection switching there is no need to coordinate the protection state between the protection controllers at both ends of the protection domain. In 1+1 bidirectional protection

switching, there is a need for a protocol to coordinate the protection state between the edges of the Protection Domain.

~~In both protection schemes traffic is restored to the working entity after the conditions causing the switchover has/have been cleared. The selection of data may revert to the traffic from the working entity if reversion is enabled, and will require coordination of protection state between the edges of the Protection Domain. To avoid frequent switching in case of intermittent defects or failures when the network is not stabilized, traffic is not switched back to the working entity before the Wait-to-Restore (WTR) timer has expired.~~

Comment [Italo Bus63]: This contradicts the paragraph just before 4.6.1.1 and it is not correct. 1+1 can be both revertive and non-revertive.

Comment [Italo Bus64]: This information is valid in general and already described in section 4.3.5.

4.6.1.1.3. P2MP Linear Protection

Linear protection may apply to protect unidirectional P2MP entity using 1+1 protection architecture. The source/root MPLS-TP node bridges the user traffic to both the working and protection entities. Each sink/leaf MPLS-TP node selects the traffic from one entity based on some predetermined criteria. Note that when there is a fault condition on one of the branches of the P2MP path, some leaf MPLS-TP nodes may select the working entity, while other leaf MPLS-TP nodes may select traffic from the protection entity.

In a 1:1 P2MP protection scheme, the source/root MPLS-TP node needs to identify the existence of a fault condition on any of the branches of the network. This requires the sink/leaf MPLS-TP nodes to notify the source/root MPLS-TP node of any fault condition. This required also a return path from the sinks/leaves to the source/root MPLS-TP node.

When protection switching is triggered, the source/root MPLS-TP node selects the protection transport path to transfer the traffic.

Comment [Italo Bus65]: Depending on the fault location, it is also possible that the root node is required to bridge the traffic to both the working and protection entities and let the leaves select.

Note that such a mechanism does not yet exist and its exact behavior is for further study.

4.6.1.2. Triggers for the Linear Protection Switching Action

The protection switching may be performed when:

- A ~~fault~~ defect condition ('failed' or 'degraded') is ~~declared~~ detected on the working entity and the protection entity has no or a lesser condition. Proactive in-band OAM CCV (Continuity and Connectivity Verification) monitoring of both the working and the protection entities may be used to enable the fast detection of a fault condition. For protection switching, it is common to run a CCV every 3.33ms. In the absence of three consecutive CCV messages,

a fault condition is declared. In order to monitor the working and the protection entities, an OAM Maintenance Entity Groups should be defined for each of the entities. OAM indications of fault conditions should be provided to the edges of the Protection Domain which are responsible for the protection switching operation. Input from OAM performance monitoring indicating degradation in the working entity may also be used as a trigger for protection switching. In the case of degradation, switching to the protection entity is needed only if the protection entity can ~~guarantee~~ exhibit better operating conditions.

- An indication is received from a lower layer server that there is a ~~network defect or~~ failure in the server layer.

- An external operator command is received (e.g., 'Forced Switch', 'Manual Switch'). For details see Section 6.1.2.

- A request to switch over is received from the far end. The far end may initiate this request for example when it gets an administrative request to switch over, or when bidirectional 1:1 protection switching is supported and there was a ~~fault~~ defect that could be detected only by the far end, etc.

As described above, the protection state should be coordinated between the end points of the Protection Domain. Control message should be exchanged between the edges of the Protection Domain to coordinate the protection state of the edge nodes. The control messages can be delivered using in-band data-plane driven control protocol or a control plane based protocol.

In order to achieve 50ms protection switching it is recommended to use in-band data-plane driven signaling protocol to coordinate the protection states. An in-band data-plane PSC (Protection State Coordination) protocol is defined in [MPLS-TP-Linear-Protection] for this purpose. This protocol is also used to detect mismatches between the configuration provisioned at the ends of the Protection Domain.

As described below in Section 6.5, GMPLS already defines procedures and messages' elements to coordinate the protection states between

Comment [Italo Bus66]: Protection triggers have been already described in section 4.1. Is there a need to describe them also here? Is it not possible to describe them in a generic section?

the edges of the protection domain. These procedures and protocols messages are specifies in [RFC4426], [RFC4872] and [RFC4873]. However, these messages lack the capability to coordinate the revertive/non-revertive behavior and the consistency of configured timers at the edges of the Protection Domain (timers such as Wait to Restore (WTR), Hold-off timer, etc.).

4.6.1.3. Applicability of Linear Protection for LSP Segments

In order to implement data-plane based linear protection on LSP segments, there is a need to support the MPLS-TP architectural element PST (Path Segment Tunnel). ~~Maintenance operations (e.g., monitoring, protection or management) engage with a transmission of messages (e.g., OAM, Protection Path Coordination, etc.) in the maintained domain. According to the MPLS architecture which is defined in [RFC3031], such messages can be initiated and terminated at the edges of a path where push and pop operations are enabled. In order to support the option to monitor, protect and manage a portion of an LSP, a new architectural element is defined, Path Segment Tunnel (PST).~~ As defined in [MPLS-TP-Framework], a Path Segment Tunnel is an LSP which is basically defined and used for the purposes of OAM monitoring, protection or management of LSP segments. PST makes use of the MPLS construct of hierarchical nested LSP which is defined in [RFC3031].

Comment [Italo Bus67]: Duplicated introduction of the PST concept. It seems better to refer to the PST definition in the MPLS-TP Framework.

For linear protection operation, PSTs should be defined over the working and protection entities between the edges of a Protection Domain. OAM and PSC messages can be initiated at the edge of the PST and sent to the peer edge of the PST. Note that these messages are sent over ~~G-ACH channels, within the PST~~the PST G-ACh and use two labels stack: ~~,~~ the PST label and, at the bottom of stack, ~~and~~ the G-~~ACH~~ ACh label (GAL).

Comment [Italo Bus68]: The G-ACh is a channel.

The end-to-end traffic of the LSP, including data-traffic and control traffic (OAM, PSC, management and signaling messages) is tunneled within the PSTs by means of label stacking as defined in [RFC3031].

The mapping between an LSP and a PST can be 1:1 which is similar to the ITU-T Tandem Connection element which defines a sub layer corresponding to a segment of a path. The mapping can also be 1:n to allow scalable protection of a set of LSPs' segments traversing the portion of the network in which a Protection Domain is defined. Note that each of these LSPs can be initiated or terminated at different endpoints in the network, but they all traverse the Protection Domain and share similar constraints (such as requirements for QoS, terms of protection ,etc.).

Note that in the context of segment protection, the PSTs serve as the working and protection entities.

Comment [HvH69]: One reviewer provided comments up to this point and ran out of review time

4.6.1.4. Shared Mesh Protection

Comment [GA70]: This is not linear protection any more, it should be numbered 4.6.2

In shared mesh protection, the protection resources are used to

protect more than one LSP that do not all have the same end points. For example, in Figure 3 there are two paths ABCDE and VWXYZ. These paths do not share end points so they cannot make use of 1:n linear protection even though they also do not share any common points of

failure.

ABCDE may be protected by the path APQRE, and VWXYZ can be protected by the path VPQRZ. In both cases, 1:1 or 1+1 protection may be used. However, it can be seen that, if 1:1 protection is used for both paths, the network segment PQR carries no traffic if there are no failures affecting either of the two working paths. Furthermore, in the event of only one failure, the segment PQR carries traffic from only one of the working paths.

Thus, it is possible for the network resources on the segment PQR to be shared by the two recovery paths. In this way, mesh protection can substantially reduce the amount of network resources that have to be reserved to provide protection of a 1:n nature.

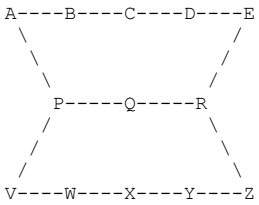

Figure 3: A Shared Mesh Protection Topology

As the complexity of the network and the number of LSPs increases, the potential for shared mesh protection also increases. However, it can rapidly become unmanageably complex. Therefore, shared mesh protection is normally pre-planned and configured by the operator, although an automated system is not out of the question.

Note that shared mesh protection operates as 1:n linear protection (see Section 4.6.1.1.1). However, the protection state needs to be ~~coordinated be~~ coordinated between a larger number of nodes: the end points of the shared concatenated protection segment (nodes P and R in the example) as well as the end points of the protected LSPs (nodes A, E, V, and Z in the example).

Additionally, note that the shared protection resources could be used to carry extra traffic. For example, in Figure 4, an LSP JPQRK could be a preemptable LSP that constitutes extra traffic over the hops PQR and would be displaced in the event of a protection event. In this case it should be noted that protection state must be additionally coordinated with the ends of the extra traffic LSPs.

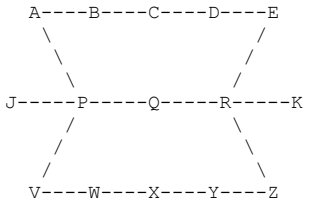

Figure 4: Shared Mesh Protection with Extra Traffic

4.6.2. Ring Networks

Several Service Providers have expresses a high level of interest in operating MPLS-TP in ring topologies and require a high level of survivability function in these topologies.

~~Different~~ Various criteria for optimization are considered in ring topologies, such as:

1. Simplification of the operation of the Ring in terms of the number of OAM Maintenance Entities that are needed to trigger the recovery actions, the number of elements of recovery, the number of management plane transactions during maintenance operations, etc.

2. Optimization of resource consumption around the ring, like the number of labels needed for the protection paths that cross the network, the total bandwidth needed in the ring to ensure the protection of the paths, etc.

[RFC5654] introduces a list of requirements on ring protection that cover the recovery mechanisms need to protect traffic in a single ring and traffic that traverses more than one ring. Note that configuration and the operation of the recovery mechanisms in a ring must scale well with the number of transport paths, the number of nodes, and the number of ring interconnects.

The requirements for ring protection are fully compatible with the generic requirements for recovery.

The architecture and the mechanisms for ring protection are specified in separate documents. These mechanisms need to be evaluated against the requirements specified in [RFC5654]. The principles for the development of the mechanisms should be:

1. Reuse existing procedures and mechanisms for recovery in ring topologies as along as their performance is as good as new potential mechanisms.

2. Ensure complete interoperability with the mechanisms defined for

arbitrary topologies to allow end-to-end protection.

4.7. Recovery in Layered Networks

In multi-layer or multi-region networking [RFC5212], recovery may be performed at multiple layers or across cascaded recovery domains.

The MPLS-TP recovery mechanism must ensure that the timing of recovery is coordinated in order to avoid races, and to allow either the recovery mechanism of the server layer to fix the problem before recovery takes place at the MPLS-TP layer, or to allow an upstream recovery domain to perform recovery before a downstream domain. In inter-connected rings, for example, it may be preferable to allow the upstream ring to perform recovery before the downstream ring, in order to ensure that recovery takes place in the ring in which the defect or failure occurred.

A hold-off timer is required to coordinate the timing of recovery at multiple layers or across cascaded recovery domains. Setting this configurable timer involves a trade-off between rapid recovery and the creation of a race condition where multiple layers respond to the same fault, potentially allocating resources in an inefficient manner. Thus, the detection of a defect or failure condition in the MPLS-TP layer should not immediately trigger the recovery process if the hold-off timer is configured to a value other than zero. Instead, the hold-off timer should be started when the defect or failure is detected and, on expiry, the recovery element should be checked to determine whether the defect or failure condition still exists. If it does exist, the defect triggers the recovery operation.

The hold-off timer should be configurable.

In other configurations, where the lower layer does not have a restoration capability, or where it is not expected to provide protection, the lower layer needs to trigger the higher layer to immediately perform recovery. Although the hold-off timer can be configured to zero to force this, it may be that with layer-independence, the higher layer does not know whether the lower layer will perform restoration. In this case, the higher layer will configure a non-zero hold-off timer and rely on a specific notification from the lower layer if the lower layer cannot perform restoration.

Reference should be made to [RFC3386] that discusses the interaction between layers in survivable networks.

4.7.1. Inherited Link-Level Protection

Where a link in the MPLS-TP network is formed from connectivity (i.e., a packet or non-packet LSP) in a lower layer network, that connectivity may itself be protected. For example, the LSP in the lower layer network may be provisioned with 1+1 protection. In this case the link in the MPLS-TP network has an inherited level of protection.

An LSP in the MPLS-TP network may be provisioned with protection in the MPLS-TP network as already described, or it may be provisioned to utilize only links that themselves have inherited protection.

By classifying the links in the MPLS-TP network according to the level of protection that they inherit from the server network, it is possible to compute an end-to-end path in the MPLS-TP network that uses only links with a specific or better level of inherited protection. This means that the end-to-end MPLS-TP LSP can be protected at the level necessary to conform with the SLA without the need to provide any additional protection in the MPLS-TP layer. This saves complexity and network resources, and reduces issues of protection switching coordination.

Where the requisite level of inherited protection is not available on all segments along the path in the MPLS-TP network, segment protection may be used to achieve the desired protection level.

It should be noted, however, that inherited protection only applies to links. Nodes cannot be protected in this way. An operator will need to perform an analysis of the relative likelihood and consequences of node failure if this approach is taken without providing any protection in the MPLS-TP LSP or PW layer to handle node failure.

4.7.2. Shared Risk Groups

When an MPLS-TP protection scheme is established, it is important that the working and protection paths do not share resources in the network. If this is not achieved, a single defect or failure may affect both the working and the protection path with the result that the traffic cannot be delivered - ~~it was, in fact,~~under such a condition the traffic is not protected.

Note that this restriction does not apply for restoration as this takes place after the fault has arisen meaning that the point of

defect or failure can be avoided if an available path exists.

When planning a recovery scheme it is possible to select paths that use diverse links and nodes within the MPLS-TP network using a topology map of the MPLS-TP layer. However, this does not guarantee that the paths are truly diverse. For example, two separate links in an MPLS-TP network may be provided by two lambdas in the same optical fiber, or by two fibers that cross the same bridge. And two completely separate MPLS-TP nodes might be situated in the same building with a shared power supply.

Thus, in order to achieve proper recovery planning, the MPLS-TP network must have an understanding of the groups of lower layer resources that share a common risk of defect or failure. From this, MPLS-TP shared risk groups can be constructed that show which MPLS-TP resources share a common risk of defect or failure. The working and protection paths can be planned to be not only node and link diverse, but to not use any resources from the same shared risk groups.

4.7.3. Fault Correlation

In a layered network a low-layer fault may be detected and reported by multiple layers and may sometimes give rise to multiple fault reports from the same layer. For example, a failure of a data link may be reported by the line cards in an MPLS-TP node, but it could also be detected and reported by the MPLS-TP OAM.

Section 4.6 explains how it is important to coordinate the survivability actions configured and operated in a multi-layer network to avoid over-equipping the survivability resources in the network, and to ensure that recovery actions are taken only in one layer at a time.

Fault correlation is about understanding what single event has led to a set of fault reports so that the recovery actions can be coordinated, and so that the fault logging system does not become overloaded. Fault correlation depends on an understanding of resource usage at lower layers, shared risk groups, and a wider view of how the layers are inter-related.

Fault correlation is most easily performed at the point of fault detection. For example, an MPLS-TP node that receives a fault notification from the lower layer and detects a fault on an LSP in the MPLS-TP layer can easily correlate these two events. Furthermore, the same node detecting multiple faults on LSPs using the same faulted data link, can easily correlate these. Such a node may use the correlation to perform group-based recovery actions, and can reduce the number of alarm events that it raises to its

management station.

Fault correlation may also be performed at a management station that receives fault reports from different layers and different nodes in the network. This enables the management station to coordinate management-originated recovery actions, and to present a consolidated fault information to the user and any automated management systems.

There is also a desire to correlate fault information detected and reported through OAM. This function would enable a fault detected at a lower layer and reported at a transit node of an MPLS-TP LSP to be correlated with an MPLS-TP layer fault detected at a Maintenance End Point (MEP) (for example the egress of the MPLS-TP LSP. Such correlation allows the coordination of recovery actions taken at the MEP, but it requires that the lower layer fault information is propagated to the MEP which is most easily achieved by using a control plane, management plane, or OAM message.

5. Applicability and Scope of Survivability in MPLS-TP

The MPLS-TP network can be viewed as two layers (the MPLS LSP layer and the PW layer). The MPLS-TP network operates over data link connections and data link networks such that the MPLS-TP links are provided by individual data links or by connections in a lower layer network. The MPLS LSP layer is a mandatory part of the MPLS-TP network, and the PW layer is an optional addition to support specific services.

MPLS-TP survivability provides recovery from defect or failure of the links and nodes in the MPLS-TP network. The link defects and failures are typically caused by defects or failures in the underlying data link connections and networks, but this section is only concerned with recovery actions taken in the MPLS-TP network, which must necessarily be to recover from the manifestation of any problem as a defect or failure in the MPLS-TP network.

This section lists which recovery elements (see Section 1) are supported in each of the two layers to recover from defects or failures of nodes or links in the MPLS-TP network.

Recovery Element	MPLS LSP Layer	PW Layer
Link Recovery	MPLS LSP recovery can be used to survive the failure of an MPLS-TP link.	The PW layer is not aware of the underlying network. This function is not supported.
Segment/Span Recovery	An individual LSP segment can be recovered to survive the failure of an MPLS-TP link.	For a SS-PW, segment recovery is the same as end-to-end recovery. Segment recovery for a MS-PW is for future study, and this function is now provided using end-to-end recovery.
Concatenated Segment Recovery	A concatenated LSP segment can be recovered to survive the failure of an MPLS-TP link or node.	Concatenated segment recovery (in a MS-PW) is for future study, and this function is now provided using end-to-end recovery.
End-to-end Recovery	An end-to-end LSP can be recovered to survive any node or link failure, except for the failure of the ingress or egress node.	End-to-end PW recovery can be applied to survive any node (including S-PE) or link failure except for the failure of the ingress egress T-PE.
Service Recovery	The MPLS LSP layer is service	PW layer service recovery requires surviving faults in

Comment [GA71]: Please expand

	agnostic. This function is not supported.	T-PEs or on ACs. This is currently out of scope for MPLS-TP.	

Comment [GA72]: Please expand

Comment [GA73]: Please expand

Table 1

Section 6 provides a description of mechanisms for survivability of MPLS-TP LSPs. Section 7 provides a brief overview of mechanisms for survivability of MPLS-TP PWs.

6. Mechanisms for Providing Survivability for MPLS-TP LSPs

This section describes the existing mechanisms available to provide protection of LSPs within MPLS-TP networks, and highlights areas where new work is required. It is expected that, as new protocol extensions and techniques are developed, this section will be updated to convert the statements of required work into references to those protocol extensions and techniques.

6.1. Management Plane

As described above, a fundamental requirement of MPLS-TP is that recovery mechanisms should be capable of functioning in the absence of a control plane. Recovery may be triggered by MPLS-TP OAM fault management functions or by external requests (e.g., an operator request for manual control of protection switching). Recovery LSPs (and in particular Restoration LSPs) may be provisioned through the management plane.

The management plane may be used to configure the recovery domain by setting the reference endpoints points (which controls the recovery actions), the working and the recovery entities, and the recovery type (e.g., 1:1 bidirectional linear protection, ring protection, etc.).

Additional parameters associated with the recovery process (such as a WTR and hold-off timers, revertive/non-revertive operation, etc.) may also be configured.

In addition, the management plane may initiate manual control of the recovery function. A priority should be set between fault conditions and operator's requests.

Since provisioning the recovery domain involves the selection of a number of options, mismatches may occur at the different reference points. The MPLS-TP OAM PSC (protection State Coordination) which is specified in [MPLS-TP-Linear-Protection] may be used as an in-band (i.e., data plane-based) control protocol to coordinate the protection states between the endpoints of the recovery domain and to check consistency of configured parameters (such as timers, revertive/non-revertive behavior, etc.) with any discovered inconsistencies being reported to the operator.

It should also be possible for the management plane to track the

recovery status by receiving reports or by issuing polls.

6.1.1. Configuration of Protection Operation

In order to implement the protection switching mechanisms, the following entities and information should be configured and provisioned:

- The endpoints of a recovery domain. As described above, these endpoints bound the element of recovery for which recovery is applied.
- The protection group which, depending on the required protection scheme, consists of a recovery entity and one or more working entities. In 1:1 or 1+1 P2P protection, in order to guarantee protection, the paths of the working entity and the recovery entities must be completely physically diverse (i.e. not share any resources or physical locations).
- As defined in Section 4.6.2, in order to implement data-plane based LSP segment recovery, there is a need to support the MPLS-TP architectural element PST (Path Segment Tunnel), since related control messages (e.g., for OAM, Protection Path Coordination, etc.) can be initiated and terminated at the edges of a path where push and pop operations are enabled. PST is an end-to-end LSP which corresponds in this context to the recovery entities (working and protection) and makes use of the MPLS construct of hierarchical nested LSP which is defined in [RFC3031]. OAM and PSC messages can be initiated at the edge of the PST and sent to the peer edge of the PST, over G-ACH. There is a need to configure the related PSTs and map between the LSP segments being protected and the PST. The mapping can be 1:1 or 1:N to allow scalable protection of a set of LSPs' segments traversing the portion of the network in which a Protection Domain is defined. Note that each of these LSPs can be initiated or terminated at different endpoints in the network, but they all traverse the Protection Domain and share similar constraints (such as requirements for QoS, terms of protection ,etc.).
- The protection type that should be defined (e.g., unidirectional 1:1, bidirectional 1+1, etc.).
- Revertive/non-revertive behavior should be configured.
- timers (such as WTR, hold-off timer, etc.) should be set.

6.1.2. External Manual Commands

The following external, manual commands may be provide for manual control of the protection switching operation. These commands apply to a protection group and they are listed in descending order of priority:

- Blocked protection action - a manual command to prevent data traffic from switching to the recovery entity. This command actually disables the protection group.

- Force protection action - a manual command that forces a switch of normal data traffic to the recovery entity.

- Manual protection action - a manual command that forces a switch of data traffic to the recovery entity when there is no defect or failure in the working or the recovery entity.

- Clear switching command - the operator may request to clear a previous administrative command to switch (manual or force switch).

Comment [HvH74]: the definition of Manual switch is wrong. The state of the working resources is not considered. The current description implies that manual switch will not take place if the working path has failed.

6.2. Fault Detection

Fault detection is a fundamental part of recovery and survivability. In all schemes except for some forms of 1+1 protection, the necessary actions for recovery of traffic delivery rely on discovering that there is some kind of fault. In 1+1 protection, the selector (at the receiving end) may simply be configured to choose the better signal, thus it does not detect a fault or degradation per se, but simply identifies which path is better delivering data.

Faults may be detected in a number of ways depending on the traffic pattern and the underlying hardware. End-to-end faults may be reported by the application or by knowledge of the application's data pattern, but this is an unusual approach. There are two more common mechanisms for detecting faults in the MPLS-TP layer:

- faults reported by the lower layers

- faults detected by protocols within the MPLS-TP layer.

In an IP/MPLS network, the second of these may utilize control plane protocols (such as the routing protocols) to detect a defect or failure of adjacency between neighboring nodes. In an MPLS-TP network, there is no certainty that a control plane will be present. Even if a control plane is present, it will be a GMPLS control plane [RFC3945] that makes a logical separation between control channels

and data channels with the result that no conclusion about the health of a data channel can be drawn from the defect or failure of an associated control channel. MPLS-TP layer faults are, therefore, only detected through the use of OAM protocols as described in Section 6.4.1.

Faults may, however, be reported by a lower layer. These generally show up as interface failures or data link failures (sometimes known as connectivity failures) within the MPLS-TP network. For example, an underlying optical link may detect loss of light and report a defect or failure of the MPLS-TP link that uses it. Alternatively, an interface card failure may be reported to the MPLS-TP layer.

Faults reported by lower layers are only visible at specific nodes within the MPLS-TP network (i.e., at the adjacent end-points of the MPLS-TP link). This would only allow recovery to be performed locally so, in order that recovery can be performed by nodes that are not immediately local to the fault, the fault must be reported (Sections 6.4.3 and 6.5.4).

6.3. Fault Isolation

If an MPLS-TP node detects that there is a fault in an LSP (that is, not a network fault reported from a lower layer, but a fault detected by examining the LSP) it can immediately perform a recovery action. However, unless the location of the fault is known, the only practical options are:

- perform end-to-end recovery
- perform some other recovery as a speculative act.

Since speculative acts are not guaranteed to achieve the desired results and could be costly, and since end-to-end recovery is a costly option, it is important to be able to isolate the fault.

Fault isolation may be achieved by dividing the network into protection domains. End-to-end protection is thereby operated on LSP segments depending on the domain in which the fault is discovered. This requires that the LSP can be monitored at the domain edges.

Alternatively, a proactive mechanism of fault isolation through OAM (Section 6.4.2) or through the control plane (Section 6.5.3) is required.

Fault isolation is particularly important for restoration because a new path must be selected that avoids the fault. It may not be

practical or desirable to select such a path that avoids the whole of the failed working path and so it is necessary to narrow down (to isolate) where the fault.

6.4. OAM Signaling

MPLS-TP provides a comprehensive set of OAM tools for fault management and performance monitoring at different nested levels (end-to-end, a portion of a path (LSP or PW) and at the link level).

These tools support proactive and on-demand fault management (for fault detection and fault localization) and for performance monitoring (to measure the quality of the signals and detect

degradation).

To support fast recovery, it is useful to use some of the proactive tools to detect fault conditions (e.g., link/node failure or degradation) and trigger the recovery action.

The MPLS-TP OAM messages run in-band with the traffic and support unidirectional and bidirectional P2P paths as well as P2MP paths.

As described in [MPLS-TP-OAM-Framework], MPLS-TP OAM operates in the context of a Maintenance Entity which bounds the OAM responsibilities and represents the portion of a path between two points which is being monitored and maintained, and in which OAM messages are exchanged. [MPLS-TP-OAM-Framework] refers also to a Maintenance Entity Group (MEG), which is a collection of one or more MEs that belongs to the same transport path (e.g., P2MP transport path) and that are maintained and monitored as a group.

An ME includes two MEPs (Maintenance Group End Points) which reside at the boundaries of an ME, and a set of zero or more MIPS (Maintenance Group Intermediate Points) which reside within the Maintenance Entity along the path. A MEP is capable of initiating and terminating OAM messages, and as such can only located at the edges of a path where push and pop operations are supported. In order to be define an ME over a portion of path there is a need to support the MPLS-TP architectural element PST (Path Segment Tunnel).

PST is an end-to-end LSP which corresponds in this context to the ME and makes use of the MPLS construct of hierarchical nested LSP which is defined in [RFC3031]. OAM messages can be initiated at the edge of the PST and sent to the peer edge of the PST, over G-ACH.

There is a need to configure the related PSTs and map between the LSP segment(s) being monitored and the PST. The mapping can be 1:1 or 1:N to allow scalable operation. Note that each of these LSPs can be

initiated or terminated at different endpoints in the network and share similar constraints (such as requirements for QoS, terms of protection ,etc.).

In the context of recovery where MPLS-TP OAM is supported, an OAM Maintenance Entity Group is defined for each of the working and protection entities.

A MIP is capable of reacting to OAM messages.

6.4.1. Fault Detection

MPLS-TP OAM tools may be used proactively to detect the following fault conditions between MEPs:

- Loss of continuity and misconnectivity - the proactive Continuity Check (CC) function is used to detect loss of continuity between two MEPs in an MEG. The proactive ~~misconnectivity~~ Connectivity verification (CV) allows a sink MEP ~~can~~ to detect a misconnectivity defect (e.g., mismerge or misconnection) with its peer source MEP when the received packet

carries an incorrect ME identifier. For protection switching, it is common to run CCV (Continuity and Connectivity Verification) message every 3.33ms. In the absence of three consecutive CCV messages, Loss of Continuity is declared and locally notified to the edge of the recovery domain to trigger a recovery action. In some cases, when a slower recovery time is acceptable, it is also possible to lengthen the transmission rate.

- Signal degradation - notification from the OAM performance monitoring indicating degradation in the working entity may also be used as a trigger for protection switching. In the case of degradation, switching to the recovery entity is needed only if the recovery entity can guarantee better conditions. Degradation can be measured by activating proactively the MPLS-TP OAM packet loss measurement or delay measurement.

- A MEP can get an indication from its sink MEP of a Remote Defect Indication and locally notify the endpoint of the recovery domain of fault condition to trigger the recovery action.

6.4.2. Testing for Faults

The management plane may be used to initiate testing of links, LSP segments, or whole LSPs.

MPLS-TP provides OAM tools which may be initiated on-demand by manual intervention for a limited time to carry out troubleshooting of links, LSP segments or whole LSPs (e.g. diagnostics, connectivity

verification, packet loss measurements, etc.). On-demand monitoring covers a combination of "in service" and "out-of service" monitoring functions. "out-of-service" testing is supported by the OAM on-demand lock operation. The lock operation temporarily disables the transport entity (LSP, LSP segment or link) for transmission of any traffic except for test traffic, and OAM (dedicated to the locked entity).

[MPLS-TP-OAM-Framework] describes the operations of the OAM functions that may be initiated on-demand and provides some considerations.

MPLS-TP supports also the in/out-of-service test operation of the recovery (protection and restoration) mechanism, the integrity of the protection/recovery transport paths and the coordination protocol between the endpoints of the recovery domain. The testing operation emulates a protection switching request without performing the actually switching action.

6.4.3. Fault Isolation

Comment [GA75]: How can OAM isolate a fault? It can be used only to locate a fault

MPLS-TP provides OAM tools to ~~isolate~~ locate a fault and determining exactly where a fault has occurred. It is often the case the fault detection only takes place at key points in the network (such as at LSP end points, or MEPs). This means that the fault may be located anywhere within a segment of the LSP concerned. Finer granularity of information is needed to implement optimal recovery actions or to diagnose the fault. On-demand tools like trace-route, loopback and on-demand CCV can be used to ~~isolate~~ locate a fault.

The information may be locally notified to the endpoint of the recovery domain to allow implementation of optimal recovery action. This may be useful in case of re-calculation of a recovery path.

The information should also be reported to the network management for diagnostics purposes.

6.4.4. Fault Reporting

The endpoints of a recovery domain should be able to report the fault conditions detected in the recovery domain to the management plane.

Comment [GA76]: Requires rewording for correct reading

In addition, a node within a recovery domain detecting a fault condition should also be able to report the fault condition to the network management. The network management should be capable to correlate the fault reports and identify the source of the fault.

MPLS-TP OAM tools support a function where an intermediate node along a path can send an alarm report message to the MEP indicating a fault condition in the server layer connecting it to its adjacent

node. The purpose of this capability is to allow a MEP to suppress alarms that may be generated as a result of the failure condition in the server layer.

6.4.5. Coordination of Recovery Actions

As described above, in some cases (such as in bidirectional protection switching, etc.) there is a need to coordinate the protection states between the edges of the recovery domain. [MPLS-TP-Linear-Protection] defines procedures and protocol messages and elements to support the PSC (Protection State Coordination) function.

The protocol is also used to signal administrative requests (e.g., manual switch, etc.) when these are provisioned only at on edge of the recovery domain.

The protocol also allow to detect mismatches between the configuration provisioned at the ends of the Protection Domain (such as timers, revertive/non-revertive behavior), and such mismatches would be reported to the management plane.

In the event that suitable coordination does not occur (because of failures of the PSC function, or because it is not run) protection switching will fail. That is, the operation of the PSC function is a fundamental part of protection switching.

6.5. Control Plane

The GMPLS control plane has been proposed as the control plane for MPLS-TP [RFC5317]. Since GMPLS was designed for use in transport networks, and has been implemented and deployed in many networks, it is not surprising that it contains many features to support a high level of survivability function.

The signaling elements of the GMPLS control plane utilize extensions to the Resource Reservation Protocol (RSVP) as documented in a series of documents commencing with [RFC3471] and [RFC3473], but based on [RFC3209] and [RFC2205]. The architecture for GMPLS is provided in [RFC3945], and [RFC4426] gives a functional description of the protocol extensions needed to support GMPLS-based recovery (i.e., protection and restoration).

A further control plane protocol called the Link Management Protocol (LMP) [RFC4204] is part of the GMPLS protocol family and can be used to coordinate fault isolation and reporting.

Clearly, the control plane techniques described here only apply where

an MPLS-TP control plane is deployed and operated. All mandatory MPLS-TP survivability features must be enabled even in the absence of the control plane, but where the control plane is present it may be used to provide alternative mechanisms ~~that may be desirable by virtue of their ease of automation or richer feature-set~~..

Comment [GA77]: The purpose of this text is unclear

6.5.1. Fault Detection

The control plane is not able to detect data plane faults. However, it does provide mechanisms to detect control plane faults and these can be used to deduce data plane faults where it is known that the control and data planes are fate sharing. Although [RFC5654] specifies that MPLS-TP must support an out-of-band control channel, it does not insist that this is used exclusively. That means that there may be deployments where an in-band (or at least in-fiber) control channel is used. In this case, the failure of the control channel can be used to infer a failure of the data channel or at least to trigger an investigation of the health of the data channel.

Both RSVP and LMP provide a control channel "keep-alive" mechanism (called the Hello message in both cases). Failure to receive a message in the configured/negotiated time period indicates a control plane failure. GMPLS routing protocols ([RFC4203] and [RFC5307] also include keepalive mechanisms designed to detect routing adjacency failures and, although these keep-alive mechanisms tend to operate at a relatively low frequency (order of seconds) it is still possible that the first indication of a control plane fault will be through the routing protocol.

Note, however, care must be taken that the failure is not caused by a problem with the control plane software or processor component at the far end of a link.

Because of the various issues involved, it is not recommended that the control plane be used as the primary mechanism for fault detection in an MPLS-TP network.

6.5.2. Testing for Faults

The control plane may be used to initiate and coordinate testing of links, LSP segments, or whole LSPs. This is important in some technologies where it is necessary to halt data transmission while testing, but may also be useful where testing needs to be

specifically enabled or configured.

LMP provides a control plane mechanism to test the continuity and connectivity (and naming) of individual links. A single management operation is required to initiate the test at one end of the link,

and LMP handles the coordination with the other end of the link. The test mechanism for an MPLS packet link relies on the LMP Test message inserted into the data stream at one end of the link and extracted at the other end of the link. This mechanism need not be disruptive to data flowing on the link.

Note that a link in LMP may in fact be an LSP tunnel used to form a link in the MPLS-TP network.

GMPLS signaling (RSVP) offers two mechanisms that may also assist with testing for faults. First, [RFC3473] defines the Admin_Status object that allows an LSP to be set into "testing mode". The interpretation of this mode is implementation specific and could be documented more precisely for MPLS-TP. The mode sets the whole LSP into a state where it can be tested; this need not be disruptive to data traffic.

The second mechanism provided by GMPLS to support testing is provided in [GMPLS-OAM]. This protocol extension supports the configuration (including enabling and disabling) of OAM mechanisms for a specific LSP.

6.5.3. Fault Isolation

Fault isolation is the process of determining exactly where a fault has occurred. It is often the case the fault detection only takes place at key points in the network (such as at LSP end points, or MEPs). This means that the fault may be located anywhere within a segment of the LSP concerned.

If segment or end-to-end protection are in use, this level of information is often sufficient to repair the LSP. However, if a finer granularity of information is needed (either to implement optimal recovery actions or to diagnose the fault), it is necessary to isolate the fault more closely.

LMP provides a cascaded test-and-propagate mechanism specifically designed for this purpose.

6.5.4. Fault Status Reporting

GMPLS signaling uses the Notify message to report fault status [RFC3473]. The Notify message can apply to a single LSP or can carry fault information for a set of LSPs to improve the scalability of fault notification.

Since the Notify message is targeted at a specific node it can be delivered rapidly without requiring hop-by-hop processing. It can be

targeted at LSP end-points, or at segment end-points (such as MEPs). The target points for Notify messages can be manually configured within the network or may be signaled as the LSP is set up. This allows the process to be made consistent with segment protection and the concept of Maintenance Entities.

GMPLS signaling also provides a slower, hop-by-hop mechanism for reporting individual LSP faults on a hop-by-hop basis using the PathErr and ResvErr messages.

[RFC4783] provides a mechanism to coordinate alarms and other event or fault information through GMPLS signaling. This mechanism is useful to understand the status of the resources used by an LSP and to help understand why an LSP is not functioning, but it is not intended to replace other fault reporting mechanisms.

GMPLS routing protocols [RFC4203] and [RFC5307] are used to advertise link availability and capabilities within a GMPLS-enabled network. Thus, the routing protocols can also provide indirect information about network faults. That is, the protocol may stop advertising or withdraw the advertisement for a failed link, or may advertise that the link is about to be shut down gracefully [GR-SHUT]. This mechanisms is, however, not normally considered to be fast enough to be used as a trigger for protection switching.

6.5.5. Coordination of Recovery Actions

Fault coordination is an important feature for certain protection mechanisms (such as bidirectional 1:1 protection). The use of the GMPLS Notify message for this purpose is described in [RFC4426], however, specific message field values remain to be defined for this operation.

A further piece of work ~~in~~ is needed from a control plane perspective to allow control and configuration of reversion behavior for end-to-end and segment protection, and the coordination of timers' values.

6.5.6. Establishment of Protection and Restoration LSPs

The management plane may be used to set up protection and recovery LSPs, but the control plane may be used if it is present.

Several protocol extensions exist to make this process more simple:

- [RFC4872] provides features in support of end-to-end protection switching.
- [RFC4873] describes how to establish a single, segment protected LSP. Note that end-to-end protection is a sub case of segment protection and [RFC4872] can be used also to provide end-to-end protection.

Internet-Draft MPLS-TP Survivability Framework March 2010

- [RFC4874] allows one LSP to be signaled with a request that its path excludes specified resources (links, nodes, SRLGs). This allows a disjoint protection path to be requested, or a recovery path to be set up avoiding failed resources.

- Lastly, it should be noted that [RFC5298] provides an overview of the GMPLS techniques available to achieve protection in multi-domain environments.

7. Pseudowire Protection Considerations

Pseudowire is one of the clients of the MPLS LSP layer of the MPLS-TP network.

It is viewed as a layer of the MPLS-TP network. Pseudowires provide end-to-end connectivity over the MPLS-TP network and may be comprised of a single pseudowire segment, or multiple segments "stitched" together to provide end-to-end connectivity.

Comment [GA78]: This contradicts the first sentence

The pseudowire may, itself, require a level of protection in order to meet the guarantees or service level of its SLA. This protection could be provided by the MPLS-TP LSPs that support the pseudowire, or could be a feature of the pseudowire layer itself.

As indicated above, the functional architecture described in this document applies to both LSPs and pseudowires. However the recovery mechanisms for pseudowires are for further study and will be defined in a separate document in the PWE3 working group.

7.1. Utilizing Underlying MPLS-TP Recovery

MPLS-TP PWs are carried across the network inside MPLS-TP LSPs. Therefore, an obvious way to protect a PW is to protect the LSP that carries it. Such protection can take any of the forms described in this document. The choice of recovery scheme will depend on the speed of recovery necessary and the traffic loss that is acceptable for the SLA that the PW is providing.

If the PW is a multi-segment PW, then LSP recovery can only protect the PW on individual segments. That is, an individual LSP recovery action cannot protect against a failure of a PW switching point (an S-PE), nor can it protect more than one segment at a time since the LSP tunnel is terminated at each S-PE. In this respect, the LSP protection of a PW is very much like the link-level protection offered to the MPLS-TP LSP layer by an underlying network layer (see Section 4.6).

7.2. Recovery in the Pseudowire Layer

Recovery in the PW layer can be provided simply by running separate PWs end-to-end. Other recovery mechanisms in the PW layer, such as segment or concatenated segment recovery, or service-level recovery involving survivability of T-PE or AC faults is for future study in a separate document.

As with any recovery mechanism, it is important to coordinate between layers. This coordination is necessary to ensure that recovery

mechanisms are only actioned in one layer at a time (that is, the recovery of an underlying LSP needs to be coordinated with the recovery of the PW itself), and to make sure that the working and protection PWs do not both use the same MPLS resources within the network (for example, by running over the same LSP tunnel - compare with Section 4.6.2).

8. Manageability Considerations

Manageability of MPLS-TP networks and function is discussed in [MPLS-TP-NM-Framework]. OAM features are discussed in [MPLS-TP-OAM-Framework].

Survivability has some key interactions with management as described in this document. In particular:

- Recovery domains may be configured such that there is not a one-to-one correspondence between the MPLS-TP network and the recovery domains.
- Survivability policies may be configured per network, per recovery domain, or per LSP.
- Configuration of OAM may involve the selection of MEPs, enabling OAM on network segments, spans, and links, and the operation of OAM on LSPs, concatenated LSP segments, and LSP segments.
- Manual commands may be used to control recovery functions including forcing recovery and locking recovery actions.

See also the consideration of security for management and OAM in Section 9 of this document

9. Security Considerations

This framework does not introduce any new security considerations, and general issues relevant to MPLS security can be found in [MPLS-SEC].

However, several points about MPLS-TP survivability should be noted here.

- If an attacker is able to force a protection switch-over, this may result in a small perturbation to user traffic, and could result in extra traffic being preempted or displaced from the protection resources. In the case of 1:n protection or shared mesh protection, it may result in other traffic becoming unprotected. Therefore, it is important that OAM protocols used to detect or notify faults use adequate security to prevent them being used (through the insertion of bogus messages, or through the capture of legitimate messages) to falsely trigger a recovery event.

- If manual commands are modified, captured, or simulated (including replay), it would be possible for an attacker to perform forced recovery actions or to impose lock-out. These actions could impact the ability to provide recovery function, and could also affect the normal operation of the network for other traffic. Therefore, management protocols used to perform manual commands must allow the operator to use appropriate security mechanisms including verification that the user issuing commands has suitable authority.

- If the control plane is used to configure or operate recovery mechanisms, the control plane protocols must also be capable of providing adequate security.

10. IANA Considerations

This informational document makes no requests for IANA action.

11. Acknowledgments

Thanks for useful comments and discussions to Italo Busi, David McWalter, Lou Berger, Yaacov Weingarten, Stewart Bryant, Dan Frost, Lievren Levrau, and Xuehui Dai.

The Editors would like to thank the participants in ITU-T Study Group 15 for their detailed review.

Some figures and text on shared mesh protection were borrowed from [MPLS-TP-MESH] with thanks to Tae-sik Cheung and Jeong-dong Ryoo.

12. References

12.1. Normative References

[RFC2205] Bradner, S., Ed., Zhang, L., Berson, S., Herzog, S., and J. Jamin, "Resource ReserVation Protocol - Version 1 Functional Specification", RFC 2205, September 1997.

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

[RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

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

[RFC4203] Kompella, K. and Y. Rekhter, "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC4204] Lang, J., Ed., "The Link Management Protocol (LMP)", RFC 4204, September 2005.

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

[RFC4428] Papadimitriou, D. and E. Mannie, "Analysis of Generalized Multi-Protocol Label Switching (GMPLS) - based Recovery Mechanisms (including Protection and Restoration) Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4428, March 2006.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, May 2007.

[RFC5307] Kompella, K. and Y. Rekhter, "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, October 2008.

[RFC5317] Bryant, S. and L. Andersson, "Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile", RFC 5317, February 2009.

[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.

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

[G.808.1] ITU-T, "Generic Protection Switching - Linear trail and subnetwork protection", Recommendation G.808.1, December 2003.

[G.841] ITU-T, "Types and Characteristics of SDH Network Protection Architectures", Recommendation G.841, October 1998.

[MPLS-TP-FWK]
Vigoureux, M., Ed., Ward, D., Ed., and Betts, "A Framework for MPLS in Transport Networks", MPLS-TP-FWK, Work in Progress.

[MPLS-TP-NM-Framework]
Mansfield, S., Gray, E., and Lam, K., "MPLS-TP Network Management Framework", draft-ietf-mpls-tp-nm-framework, Work in Progress.

[MPLS-TP-OAM]
Buci, I., Ed. and B. Niven-Jenkins, Ed., "Requirements for OAM in MPLS Transport Networks", draft-ietf-mpls-tp-oam-requirements, Work in Progress.

[MPLS-TP-OAM-Framework]
Buci, I., Ed. and B. Niven-Jenkins, Ed., "A Framework for MPLS in Transport Networks", draft-ietf-mpls-tp-oam-framework, Work in Progress.

12.2. Informative References

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3386] Lai, W. and D. McDysan, "Network Hierarchy and Multilayer Survivability", RFC 3386, November 2002.

[RFC3469] Sharma, V. and F. Hellstrand, "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, February 2003.

[RFC4426] Lang, J., Ed., Rajagopalan, B., and D. Papadimitriou, "Generalized Multiprotocol Label Switching (GMPLS) Recovery Functional Specification", RFC 4426, March 2006.

[RFC4726] Farrel, A., Vasseur, J.-P., and Ayyangar, A., "A Framework for Inter-Domain Multiprotocol Label Switching Traffic Engineering", RFC 4726, November 2006.

[RFC4783] Berger, L., "GMPLS - Communication of Alarm Information", RFC 4783, December 2006.

[RFC4872] Lang, J., Rekhter, Y., and D. Papadimitriou, "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.

[RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol- Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

[RFC5212] Shiomoto, K., Papadimitriou, D., Le Roux, JL., Vigoureux, M., and Brungard, D., " Requirements for GMPLS-Based Multi-Region and Multi-Layer Networks (MRN/MLN)", RFC 5212, July 2008

[RFC5298] Takeda, T., Farrel, A., Ikejiri, Y., and JP. Vasseur, "Analysis of Inter-Domain Label Switched Path (LSP) Recovery", RFC 5298, August 2008.

[G.8081] ITU-T, "Terms and definitions for Automatically Switched Optical Networks (ASON)", Recommendation G.8081, June 2004 and Recommendation G.8081 Amendment 1, June 2006.

[GMPLS-OAM]
Takacs, A., Fedyk, D., and H. Jia, "OAM Configuration Framework and Requirements for GMPLS RSVP-TE", draft-ietf-ccamp-oam-configuration-fwk, Work in Progress.

[GR-SHUT] Ali, Z., Vasseur, J.-P., Zamfir, A., and Newton, J., "Graceful Shutdown in MPLS and Generalized MPLS Traffic Engineering Networks", draft-ietf-ccamp-mpls-graceful-shutdown, Work in Progress.

[MPLS-SEC] L. Fang (Ed.), " Security Framework for MPLS and GMPLS Networks", draft-ietf-mpls-mpls-and-gmpls-security-framework, Work in Progress.

[MPLS-TP-Linear-Protection] Weingarten, Y., Bryant, S., Ed., Sprecher, N., Ed., Van Helvoort, H., Ed., and A. Fulignoli, "MPLS-TP Linear Protection", draft-ietf-mpls-tp-linear-protection, Work in Progress.

[MPLS-TP-MESH] Cheung , T., and Ryoo, J., "MPLS-TP Mesh Protection", draft-cheung-mpls-tp-mesh-protection, Work in Progress.

[OAM-SOUP] Andersson, L., Betts, M., Van Helvoort, H., Bonica, R., and D. Romascanu, "MPLS-TP Linear Protection", draft-ietf-opsawg-mpls-tp-oam-def, Work in Progress.

[ROSETTA] Van Helvoort, H., Ed., Andersson, L., and N. Sprecher, "A Thesaurus for the Terminology used in Multiprotocol Label Switching Transport Profile (MPLS-TP) drafts/RFCs and ITU-T's Transport Network Recommendations", draft-ietf-mpls-tp-rosetta-stone, Work in Progress.

Authors' Addresses

Nurit Sprecher
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
Hod Hasharon, 45241
Israel

Email: nurit.sprecher@nsn.com

Adrian Farrel
Old Dog Consulting

Email: adrian@olddog.co.uk