

IETF Draft
Multi-Protocol Label Switching
Expires: January 2001

Srinivas Makam
Vishal Sharma
Ken Owens
Changcheng Huang
Tellabs Operations, Inc.

Fiffi Hellstrand
Jon Weil
Loa Andersson
Bilel Jamoussi
Nortel Networks

Brad Cain
Mirror Image Internet

Seyhan Civanlar
Coreon Networks

Angela Chiu
AT&T Labs

July 2000

Framework for MPLS-based Recovery

<[draft-makam-mpls-recovery-frmrk-01.txt](#)>

Status of this memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

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>.

Abstract

Multi-protocol label switching (MPLS) [1] integrates the label swapping forwarding paradigm with network layer routing. To deliver reliable service, MPLS requires a set of procedures to provide protection of the traffic carried on different paths. This requires that the label switched routers (LSRs) support fault detection, fault notification, and fault recovery mechanisms, and that MPLS signaling [2] [3] [4] [5] [6] support the configuration of recovery. With these objectives in mind, this document specifies a framework for MPLS based recovery.

Table of Contents	Page
1.0 Introduction	3
1.1 Background	3
1.2 Motivations for MPLS-Based Recovery	4
1.3 Objectives	5
 2.0 Overview	 6
2.1 Recovery Models	6
2.2 Recovery Cycles	8
2.2.1 MPLS Recovery Cycle Model	8
2.2.2 MPLS Reversion Cycle Model	10
2.2.3 Dynamic Reroute Cycle Model	11
2.3 Terminology	13
2.4 Abbreviations	17
 3.0 MPLS Recovery Principles	 17
3.1 Configuration of Recovery	17
3.2 Initiation of Path Setup	18
3.3 Initiation of Resource Allocation	18
3.4 Scope of Recovery	19
3.4.1 Topology	19
3.4.1.1 Local Repair	19
3.4.1.2 Global Repair	20
3.4.1.3 Alternate Egress Repair	20
3.4.1.4 Multi-Layer Repair	21
3.4.1.5 Concatenated Protection Domains	21
3.4.2 Path Mapping	21
3.4.3 Bypass Tunnels	22
3.4.4 Recovery Granularity	23
3.4.4.1 Selective Traffic Recovery	23
3.4.4.2 Bundling	23
3.4.5 Recovery Path Resource Use	23
3.5 Fault Detection	24
3.6 Fault Notification	25

3.7	Switch Over Operation	25
3.7.1	Recovery Trigger	25
3.7.2	Recovery Action	26
3.8	Switch Back Operation	26
3.8.1	Revertive and Non-revertive Mode	26
3.8.2	Restoration and Notification	27
3.8.3	Reverting to Preferred LSP	28
3.9	Performance	28
4.0	Recovery Requirements	28
5.0	MPLS Recovery Options	29
6.0	Comparison Criteria	30
7.0	Security Considerations	32
8.0	Intellectual Property Considerations	32
9.0	Acknowledgements	32
10.0	Author's Addresses	33
11.0	References	34

[1.0](#) Introduction

This memo describes a framework for MPLS-based recovery. We provide a detailed taxonomy of recovery terminology, and discuss the motivation for, the objectives of, and the requirements for MPLS-based recovery. We outline principles for MPLS-based recovery, and also provide comparison criteria that may serve as a basis for comparing and evaluating different recovery schemes.

[1.1](#) Background

Network routing deployed today is focussed primarily on connectivity and typically supports only one class of service, the best effort class. Multi-protocol label switching, on the other hand, by integrating forwarding based on label-swapping of a link local label with network layer routing allows flexibility in the delivery of new routing services. MPLS allows for using media specific forwarding mechanisms as label swapping. This enables more sophisticated features such as quality-of-service (QoS) and traffic engineering [7] to be implemented more effectively. An important component of providing QoS, however, is the ability to transport data reliably and efficiently. Although the current routing algorithms are very robust and survivable, the amount of time they take to recover from a fault can be significant, on the order of several seconds or minutes, causing serious disruption of service for some applications in the interim. This is unacceptable to many organizations that aim to provide a highly reliable service, and thus require recovery

times on the order of tens of milliseconds, as specified, for example, in the GR253 specification for SONET.

Since MPLS is likely to be the technology of choice in the future IP-based transport network, it is imperative that MPLS be able to provide protection and restoration of traffic. In fact, a protection priority could be used as a differentiating mechanism for premium services that require high reliability. The remainder of this document provides a framework for MPLS based recovery. It is focused at a conceptual level and is meant to address motivation, objectives and requirements. Issues of mechanism, policy, routing plans and characteristics of traffic carried by protection paths are beyond the scope of this document.

1.2 Motivation for MPLS-Based Recovery

MPLS based protection of traffic (called MPLS-based Recovery) is useful for a number of reasons. The most important is its ability to increase network reliability by enabling a faster response to faults than is possible with traditional Layer 3 (or the IP layer) alone while still providing the visibility of the network afforded Layer 3. Furthermore, a protection mechanism using MPLS could enable IP traffic to be put directly over WDM optical channels, without an intervening SONET layer. This would facilitate the construction of IP-over-WDM networks.

The need for MPLS-based recovery arises because of the following:

- I. Layer 3 or IP rerouting may be too slow for a core MPLS network that needs to support high reliability/availability.
- II. Layer 0 (for example, optical layer) or Layer 1 (for example, SONET) mechanisms may be deployed in ring topologies and may not always include mesh protection. That is, layer 0 or layer 1 networks may not be deployed in topologies that meet carriers' protection goals.
- III. The granularity at which the lower layers may be able to protect traffic may be too coarse for traffic that is switched using MPLS-based mechanisms.
- IV. Layer 0 or Layer 1 mechanisms may have no visibility into higher layer operations. Thus, while they may provide, for example, link protection, they cannot easily provide node protection.

Furthermore there is a need for open standards.

V. Establishing interoperability of protection mechanisms between routers/LSRs from different vendors in IP or MPLS networks is urgently required to enable the adoption of MPLS as a viable core transport and traffic engineering technology.

1.3 Objectives/Goals

We lay down the following objectives for MPLS-based recovery.

I. MPLS-based recovery mechanisms should facilitate fast (10's of ms) recovery times.

II. MPLS-based recovery should maximize network reliability and availability. MPLS based protection of traffic should minimize the number of single points of failure in the MPLS protected domain.

III. MPLS-based recovery techniques should be applicable for protection of traffic at various granularities. For example, it should be possible to specify MPLS-based recovery for a portion of the traffic on an individual path, for all traffic on an individual path, or for all traffic on a group of paths.

IV. MPLS-based recovery techniques may be applicable for an entire end-to-end path or for segments of an end-to-end path.

V. MPLS-based recovery actions should not adversely affect other network operations.

VI. MPLS-based recovery actions in one MPLS protection domain (defined in [Section 2.2](#)) should not adversely affect the recovery actions in other MPLS protection domains.

VII. MPLS-based recovery mechanisms should be able to take into consideration the recovery actions of lower layers.

VIII. MPLS-based recovery actions should avoid network-layering violations. That is, defects in MPLS-based mechanisms should not trigger lower layer protection switching.

IX. MPLS-based recovery mechanisms should minimize the loss of data and packet reordering during recovery operations. (The current MPLS specification has itself no explicit requirement on reordering).

X. MPLS-based recovery mechanisms should minimize the state overhead incurred for each recovery path maintained.

XI. MPLS-based recovery mechanisms should be able to preserve the constraints on traffic after switchover, if desired. That is, if

desired, the recovery path should meet the resource requirements of, and achieve the same performance characteristics, as the working path.

2.0 Overview

There are several options for providing protection of traffic using MPLS. The most generic requirement is the specification of whether recovery should be via Layer 3 (or IP) rerouting or via MPLS protection switching or rerouting actions.

Generally network operators aim to provide the fastest and the best protection mechanism that can be provided at a reasonable cost. The higher the level of protection, the more resources it consumes. MPLS-based recovery should give the flexibility to select the recovery mechanism, choose the granularity at which traffic is protected, and to also choose the specific types of traffic that are protected in order to give operators more control over that tradeoff. With MPLS-based recovery, it can be possible to provide different levels of protection for different classes of service, based on their service requirements. For example, using approaches outlined below, a VLL service that supports real-time applications like VoIP may be supported using link/node protection together with pre-established, pre-reserved path protection, while best effort traffic may use established-on-demand path protection or simply rely on IP re-route or higher layer recovery mechanisms. As another example of their range of application, MPLS-based recovery strategies may be used to protect traffic not originally flowing on label switched paths, such as IP traffic that is normally routed hop-by-hop, as well as traffic forwarded on label switched paths.

2.1 Recovery Models

There are two basic models for path recovery: rerouting and protection switching.

Protection switching and rerouting, as defined below, may be used together. For example, protection switching to a recovery path may be used for rapid restoration of connectivity while rerouting determines a new optimal network configuration, rearranging paths, as needed, at a later time [8] [9].

2.1.1 Rerouting

Recovery by rerouting is defined as establishing new paths or path segments on demand for restoring traffic after the occurrence of a

Makam, et al.

fault. The new paths may be based upon fault information, network routing policies, pre-defined configurations and network topology information. Thus, upon detecting a fault, the affected paths are re-established using signaling. Reroute mechanisms are inherently slower than protection switching mechanisms, since more must be done following the detection of a fault. Once the network routing algorithms have converged after a fault, it may be preferable, in some cases, to reoptimize the network by performing a reroute based on the current state of the network and network policies. This is currently discussed further in [Section 3.8](#), but will also be clarified further in upcoming revisions of this document.

In terms of the principles defined in [section 3](#), reroute recovery employs paths established-on-demand with resources reserved-on-demand.

[2.1.2](#) Protection Switching

Protection switching recovery mechanisms pre-establish a recovery path or path segment, based upon network routing policies, the restoration requirements of the traffic on the working path, and administrative considerations. The recovery path may or may not be link and node disjoint with the working path [[10](#)]. When a fault is detected, the affected traffic that is considered for protection is switched over to the recovery path(s) and restored.

In terms of the principles in [section 3](#), protection switching employs pre-established recovery paths, and if resource reservation is required on the recovery path, pre-reserved resources.

[2.1.2.1](#). Subtypes of Protection Switching

The resources (bandwidth, buffers, processing) on the recovery path may be used to carry either a copy of the working path traffic or extra traffic that is displaced when a protection switch occurs. This leads to two subtypes of protection switching.

In 1+1 ("one plus one") protection, the resources (bandwidth, buffers, processing capacity) on the recovery path are fully reserved, if needed, and carry the same traffic as the working path. Selection between the traffic on the working and recovery paths is made at the path merge LSR (PML).

In 1:1 ("one for one") protection, the resources (if any) allocated on the recovery path are fully available to preemptible low priority traffic except when the recovery path is in use due to a fault on the working path. In other words, in 1:1 protection, the protected traffic normally travels only on the working path, and is switched

to the recovery path only when the working path has a fault. Once the protection switch is initiated, the low priority traffic being carried on the recovery path may be displaced by the protected traffic. This method affords a way to make efficient use of the recovery path resources.

This concept can be extended to 1:n (one for n) and m:n (m for n) protection.

Additional specifications of the recovery actions are found in [Section 3](#).

[2.2](#) The Recovery Cycles

There are three defined recovery cycles; the MPLS Recovery Cycle, the MPLS Reversion Cycle and the Dynamic Re-routing Cycle. The first cycle detects a fault and restores traffic onto MPLS-based recovery paths. If the recovery path is non-optimal the cycle may be followed by any of the two latter to achieve an optimized network again. The reversion cycle applies for explicitly routed traffic that that does not rely on any dynamic routing protocols to be converged. The dynamic re-routing cycle applies for traffic that is forwarded based on hop-by-hop routing.

[2.2.1](#) MPLS Recovery Cycle Model

The MPLS recovery cycle model is illustrated in Figure 1. Definitions and a key to abbreviations follow.

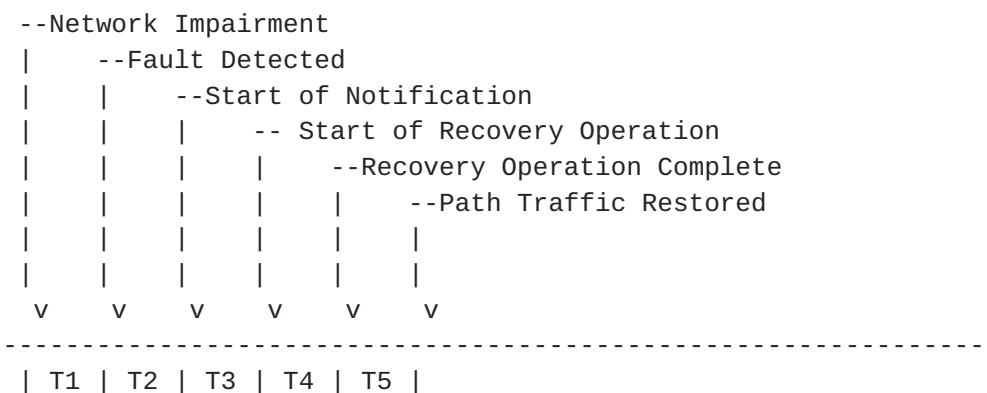


Figure 1. MPLS Recovery Cycle Model

The various timing measures used in the model are described below.

T1 Fault Detection Time

T2 Hold-off Time

- T3 Notification Time
- T4 Recovery Operation Time
- T5 Traffic Restoration Time

Definitions of the recovery cycle times are as follows:

Fault Detection Time

The time between the occurrence of a network impairment and the moment the fault is detected by MPLS-based recovery mechanisms. This time may be highly dependent on lower layer protocols.

Hold-Off Time

The configured waiting time between the detection of a fault and taking MPLS-based recovery action, to allow time for lower layer protection to take effect. The Hold-off Time may be zero.

Note: The Hold-Off Time may occur after the Notification Time interval if the node responsible for the switchover, the Path Switch LSR (PSL), rather than the detecting LSR, is configured to wait.

Notification Time

The time between initiation of an FIS by the LSR detecting the fault and the time at which the Path Switch LSR (PSL) begins the recovery operation. This is zero if the PSL detects the fault itself.

Note: If the PSL detects the fault itself, there still may be a Hold-Off Time period between detection and the start of the recovery operation.

Recovery Operation Time

The time between the first and last recovery actions. This may include message exchanges between the PSL and PML to coordinate recovery actions.

Traffic Restoration Time

The time between the last recovery action and the time that the traffic (if present) is completely - recovered. This interval is intended to account for the time required for traffic to once again arrive at the point in the network that experienced disrupted or degraded service due to the occurrence of the fault (e.g. the PML). This time may depend on the location of the fault, the recovery mechanism, and the propagation delay along the recovery path.

2.2.2 MPLS Reversion Cycle Model

Protection switching, revertive mode, requires the traffic to be switched back to a preferred path when the fault on that path is cleared. The MPLS reversion cycle model is illustrated in Figure 2. Note that the cycle shown below comes after the recovery cycle shown in Fig. 1.

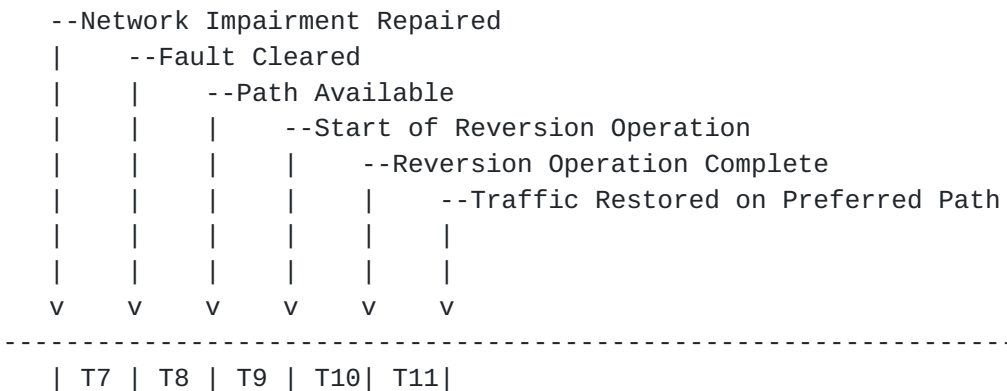


Figure 2. MPLS Reversion Cycle Model

The various timing measures used in the model are described below.

T7 Fault Clearing Time
T8 Wait-to-Restore Time
T9 Notification Time
T10 Reversion Operation Time
T11 Traffic Restoration Time

Note that time T6 (not shown above) is the time for which the network impairment is not repaired and traffic is flowing on the recovery path.

Definitions of the reversion cycle times are as follows:

Fault Clearing Time

The time between the repair of a network impairment and the time that MPLS-based mechanisms learn that the fault has been cleared. This time may be highly dependent on lower layer protocols.

Wait-to-Restore Time

The configured waiting time between the clearing of a fault and MPLS-based recovery action(s). Waiting time may be needed to ensure the path is stable and to avoid flapping in cases where a fault is

Makam, et al. [Page 10]

intermittent. The Wait-to-Restore Time may be zero.

Note: The Wait-to-Restore Time may occur after the Notification Time interval if the PSL is configured to wait.

Notification Time

The time between initiation of an FRS by the LSR clearing the fault and the time at which the path switch LSR begins the reversion operation. This is zero if the PSL clears the fault itself.

Note: If the PSL clears the fault itself, there still may be a Wait-to-Restore Time period between fault clearing and the start of the reversion operation.

Reversion Operation Time

The time between the first and last reversion actions. This may include message exchanges between the PSL and PML to coordinate reversion actions.

Traffic Restoration Time

The time between the last reversion action and the time that traffic (if present) is completely restored on the preferred path. This interval is expected to be quite small since both paths are working and care may be taken to limit the traffic disruption (e.g., using "make before break" techniques and synchronous switch-over).

In practice, the only interesting times in the reversion cycle are the Wait-to-Restore Time and the Traffic Restoration Time (or some other measure of traffic disruption). Given that both paths are available, there is no need for rapid operation, and a well-controlled switch-back with minimal disruption is desirable.

2.2.3 Dynamic Re-routing Cycle Model

Dynamic rerouting aims to bring the IP network to a stable state after a network impairment has occurred. A re-optimized network is achieved after the routing protocols have converged, and the traffic is moved from a recovery path to a (possibly) new working path. The steps involved in this mode are illustrated in Figure 3.

Note that the cycle shown below may follow the recovery cycle shown in Fig. 1 or the reversion cycle shown in Fig. 2, or both (in the event that both the recovery cycle and the reversion cycle take place before the routing protocols converge, and after the convergence of the routing protocols it is determined (based on on-Makam, et al.

line algorithms or off-line traffic engineering tools, network configuration, or a variety of other possible criteria) that there is a better route for the working path).

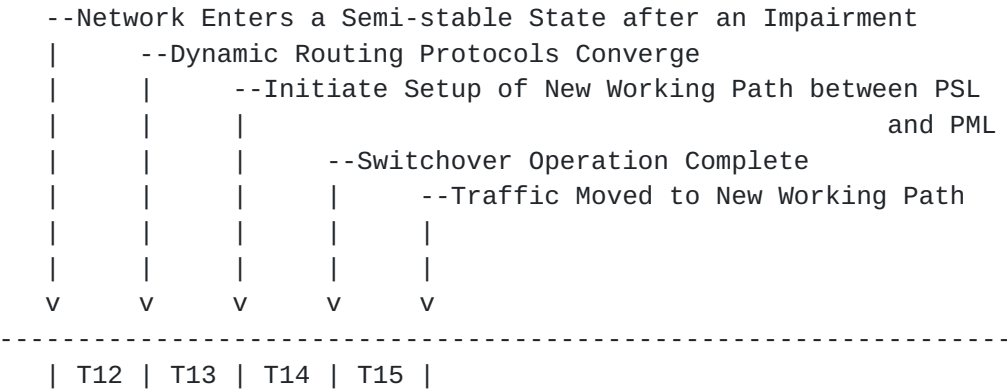


Figure 3. Dynamic Rerouting Cycle Model

The various timing measures used in the model are described below.

- T12 Network Route Convergence Time
- T13 Hold-down Time (optional)
- T14 Switchover Operation Time
- T15 Traffic Restoration Time

Network Route Convergence Time

We define the network route convergence time as the time taken for the network routing protocols to converge and for the network to reach a stable state.

Holddown Time

We define the holddown period as a bounded time for which a recovery path must be used. In some scenarios it may be difficult to determine if the working path is stable. In these cases a holddown time may be used to prevent excess flapping of traffic between a working and a recovery path.

Switchover Operation Time

The time between the first and last switchover actions. This may include message exchanges between the PSL and PML to coordinate the switchover actions.

As an example of the recovery cycle, we present a sequence of events that occur after a network impairment occurs and when a protection switch is followed by dynamic rerouting.

- I. Link or path fault occurs
- II. Signaling initiated (FIS) for the fault detected
- III. FIS arrives at the PSL
- IV. The PSL initiates a protection switch to a pre-configured recovery path
- V. The PSL switches over the traffic from the working path to the recovery path
- VI. The network enters a semi-stable state
- VII. Dynamic routing protocols converge after the fault, and a new working path is calculated (based, for example, on some of the criteria mentioned earlier in [Section 2.1.1](#)).
- VIII. A new working path is established between the PSL and the PML (assumption is that PSL and PML have not changed)
- IX. Traffic is switched over to the new working path.

[2.3](#) Definitions and Terminology

This document assumes the terminology given in [[11](#)], and, in addition, introduces the following new terms.

[2.3.1](#) General Recovery Terminology

Rerouting

A recovery mechanism in which the recovery path or path segments are created dynamically after the detection of a fault on the working path. In other words, a recovery mechanism in which the recovery path is not pre-established.

Protection Switching

A recovery mechanism in which the recovery path or path segments are created prior to the detection of a fault on the working path. In other words, a recovery mechanism in which the recovery path is pre-established.

Working Path

The protected path that carries traffic before the occurrence of a fault. The working path exists between a PSL and PML. The working path can be of different kinds; a hop-by-hop routed path, a trunk, a Makam, et al.

link, an LSP or part of a multipoint-to-point LSP.
Two synonyms for a working path are primary path, active path.

Recovery Path

The path by which traffic is restored after the occurrence of a fault. In other words, the path on which the traffic is directed by the recovery mechanism. The recovery path is established by MPLS means. The recovery path can either be an equivalent recovery path and ensure no reduction in quality of service, or be a limited recovery path and thereby not guarantee the same quality of service (or some other criteria of performance) as the working path. A limited recovery path is not expected to be used for an extended period of time.
Synonyms for a recovery path are; back-up path, alternative path, protection path.

Path Group (PG)

A logical bundling of multiple working paths, each of which is routed identically between a Path Switch LSR and a Path Merge LSR.

Protected Path Group (PPG)

A path group that requires protection.

Protected Traffic Portion (PTP)

The portion of the traffic on an individual path that requires protection. For example, code points in the EXP bits of the shim header may identify a protected portion.

Path Switch LSR (PSL)

An LSR that is the transmitter of both the working path traffic and its corresponding recovery path traffic. The PSL is responsible for switching of the traffic between the working path and the recovery path.

Path Merge LSR (PML)

An LSR that receives both working path traffic and its corresponding recovery path traffic, and either merges their traffic into a single outgoing path, or, if it is itself the destination, passes the traffic on to the higher layer protocols.

Intermediate LSR

An LSR on a working or recovery path that is neither a PSL nor a PML for that path.

Bypass Tunnel

A path that serves to backup a set of working paths using the label stacking approach. The working paths and the bypass tunnel must all share the same path switch LSR (PSL) and the path merge LSR (PML).

Switch-Over

The process of switching the traffic from the path that the traffic is flowing on onto one or more alternate path(s). This may involve moving traffic from a working path onto one or more recovery paths, or may involve moving traffic from a recovery path(s) on to a more optimal working path(s).

Switch-Back

The process of returning the traffic from one or more recovery paths back to the working path(s).

Revertive Mode

A recovery mode in which traffic is automatically switched back from the recovery path to the original working path upon the restoration of the working path to a fault-free condition.

Non-revertive Mode

A recovery mode in which traffic is not automatically switched back to the original working path after this path is restored to a fault-free condition. (Depending on the configuration, the original working path may, upon moving to a fault-free condition, become the recovery path, or it may be used for new working traffic, and be no longer associated with its original recovery path).

MPLS Protection Domain

The set of LSRs over which a working path and its corresponding recovery path are routed.

MPLS Protection Plan

The set of all LSP protection paths and the mapping from working to protection paths deployed in an MPLS protection domain at a given time.

Liveness Message

A message exchanged periodically between two adjacent LSRs that serves as a link probing mechanism. It provides an integrity check of the forward and the backward directions of the link between the two LSRs as well as a check of neighbor aliveness.

Path Continuity Test

A test that verifies the integrity and continuity of a path or path segment. The details of such a test are beyond the scope of this draft. (This could be accomplished, for example, by transmitting a control message along the same links and nodes as the data traffic.)

2.3.2 Failure Terminology

Path Failure (PF)

Path failure is fault detected by MPLS-based recovery mechanisms, which is define as the failure of the liveness message test or a path continuity test, which indicates that path connectivity is lost.

Path Degraded (PD)

Path degraded is a fault detected by MPLS-based recovery mechanisms that indicates that the quality of the path is unacceptable.

Link Failure (LF)

A lower layer fault indicating that link continuity is lost. This may be communicated to the MPLS-based recovery mechanisms by the lower layer.

Link Degraded (LD)

A lower layer indication to MPLS-based recovery mechanisms that the link is performing below an acceptable level.

Fault Indication Signal (FIS)

A signal that indicates that a fault along a path has occurred. It is relayed by each intermediate LSR to its upstream or downstream neighbor, until it reaches an LSR that is setup to perform MPLS recovery.

Fault Recovery Signal (FRS)

A signal that indicates a fault along a working path has been repaired. Again, like the FIS, it is relayed by each intermediate

LSR to its upstream or downstream neighbor, until it reaches the LSR that performs recovery of the original path.

2.4 Abbreviations

FIS: Fault Indication Signal.
FRS: Fault Recovery Signal.
LD: Link Degraded.
LF: Link Failure.
PD: Path Degraded.
PF: Path Failure.
PML: Path Merge LSR.
PG: Path Group.
PPG: Protected Path Group.
PTP: Protected Traffic Portion.
PSL: Path Switch LSR.

3.0 MPLS-based Recovery Principles

MPLS-based recovery refers to the ability to effect quick and complete restoration of traffic affected by a fault in an MPLS-enabled network. The fault may be detected on the IP layer or in lower layers over which IP traffic is transported. Fast MPLS protection may be viewed as the MPLS LSR switch completion time that is comparable to, or equivalent to, the 50 ms switch-over completion time of the SONET layer. This section provides a discussion of the concepts and principles of MPLS-based recovery. The concepts are presented in terms of atomic or primitive terms that may be combined to specify recovery approaches. We do not make any assumptions about the underlying layer 1 or layer 2 transport mechanisms or their recovery mechanisms.

3.1 Configuration of Recovery

An LSR should allow for configuration of the following recovery options:

Default-recovery (No MPLS-based recovery enabled): Traffic on the working path is recovered only via Layer 3 or IP rerouting. This is equivalent to having no MPLS-based recovery. This option may be used for low priority traffic or for traffic that is recovered in another way (for example load shared traffic on parallel working paths may be automatically recovered upon a fault along one of the working paths by distributing it among the remaining working paths)

Recoverable (MPLS-based recovery enabled): This working path is recovered using one or more recovery paths, either via rerouting or via protection switching.

3.2 Initiation of Path Setup

There are three options for the initiation of the recovery path setup.

Pre-established:

This is the same as the protection switching option. Here a recovery path(s) is established prior to any failure on the working path. The path selection can either be determined by an administrative centralized tool (online or offline), or chosen based on some algorithm implemented at the PSL and possibly intermediate nodes. To guard against the situation when the pre-established recovery path fails before or at the same time as the working path, the recovery path should have secondary configuration options as explained in [Section 3.3](#) below.

Pre Qualified:

A pre-established path need not be created, it may be pre-qualified. A pre-qualified recovery path is not created expressly for protecting the working path, but instead is a path created for other purposes that is designated as a recovery path after determination that it is an acceptable alternative for carrying the working path traffic.

Established-on-Demand:

This is the same as the rerouting option. Here, a recovery path is established after a failure on its working path has been detected and notified to the PSL.

Additional options are possible as MPLS is extended to control optical networks. One example of this is shared mesh protection in optical networks where the wavelength (or port) in-to-out mapping for a recovery lightpath is selected in every optical layer cross-connect prior to the failure, but the physical cross-connect is not made until after the failure occurs. This and other options related to optical MPLS are for further study.

3.3 Initiation of Resource Allocation

A recovery path may support the same traffic contract as the working path, or it may not. We will distinguish these two situations by
Makam, et al. [Page 18]

using different additive terms. If the recovery path is capable of replacing the working path without degrading service, it will be called an equivalent recovery path. If the recovery path lacks the resources (or resource reservations) to replace the working path without degrading service, it will be called a limited recovery path. Based on this, there are two options for the initiation of resource allocation:

Pre-reserved:

This option applies only to protection switching. Here a pre-established recovery path reserves required resources on all hops along its route during its establishment. Although the reserved resources (e.g., bandwidth and/or buffers) at each node cannot be used to admit more working paths, they are available to be used by all traffic that is present at the node before a failure occurs, which results in better resource usage than SONET APS.

Reserved-on-Demand:

This option may apply either to rerouting or to protection switching. Here a recovery path reserves the required resources after a failure on the working path has been detected and notified to the PSL and before the traffic on the working path is switched over to the recovery path.

Note that under both the options above, depending on the amount of resources reserved on the recovery path, it could either be an equivalent recovery path or a limited recovery path.

3.4 Scope of Recovery

3.4.1 Topology

3.4.1.1 Local Repair

The intent of local repair is to protect against a single link or neighbor node fault. In local repair (also known as local recovery [12] [9]), the node detecting the fault is the one to initiate recovery (either rerouting or protection switching). Local repair can be of two types:

Link Recovery/Restoration

In this case, the recovery path may be configured to route around a certain link deemed to be unreliable. If protection switching is used, several recovery paths may be configured for one working path, depending on the specific faulty link that each protects against.

Alternatively, if rerouting is used, upon the occurrence of a fault on the specified link each path is rebuilt such that it detours around the faulty link.

In this case, the recovery path need only be disjoint from its working path at a particular link on the working path, and may have overlapping segments with the working path. Traffic on the working path is switched over to an alternate path at the upstream LSR that connects to the failed link. This method is potentially the fastest to perform the switchover, and can be effective in situations where certain path components are much more unreliable than others.

Node Recovery/Restoration

In this case, the recovery path may be configured to route around a neighbor node deemed to be unreliable. Thus the recovery path is disjoint from the working path only at a particular node and at links associated with the working path at that node. Once again, the traffic on the primary path is switched over to the recovery path at the upstream LSR that directly connects to the failed node, and the recovery path shares overlapping portions with the working path.

3.4.1.2 Global Repair

The intent of global repair is to protect against any link or node fault on the entire path or on a segment of a path (with the obvious exception of the ingress and egress nodes). In global repair (also known as path recovery/restoration) the node that initiates the recovery may be distant from the faulty link or node. In some cases, a fault notification (in the form of a FIS) must be sent from the node detecting the fault to the PSL. In many cases, the recovery path can be made completely link and node disjoint with its working path. This has the advantage of protecting against all link and node fault(s) on the working path (or path segment), and being more efficient than per-hop link or node recovery.

In addition, it can be potentially more optimal in resource usage than the link or node recovery. However, it is in some cases slower than local repair since it takes longer for the fault notification message to get to the PSL to trigger the recovery action.

3.4.1.3 Alternate Egress Repair

It is possible to restore service without specifically recovering the faulted path.

For example, for best effort IP service it is possible to select a recovery path that has a different egress point from the working

path (i.e., there is no PML). The recovery path egress must simply be a router that is acceptable for forwarding the FEC carried by the working path (without creating looping). In an engineering context, specific alternative FEC/LSP mappings with alternate egresses can be formed.

3.4.1.4 Multi-Layer Repair

Multi-layer repair broadens the network designer's tool set for those cases where multiple network layers can be managed together to achieve overall network goals. Specific criteria for determining when multi-layer repair is appropriate are beyond the scope of this draft.

3.4.1.5 Concatenated Protection Domains

A given service may cross multiple networks and these may employ different recovery mechanisms. It is possible to concatenate protection domains so that service recovery can be provided end-to-end. It is considered that the recovery mechanisms in different domains may operate autonomously, and that multiple points of attachment may be used between domains (to ensure there is no single point of failure). Details of concatenated protection domains are beyond the scope of this draft.

3.4.2 Path Mapping

Path mapping refers to the methods of mapping traffic from a faulty working path on to the recovery path. There are several options for this, as described below. Note that the options below should be viewed as atomic terms that only describe how the working and protection paths are mapped to each other. The issues of resource reservation along these paths, and how switchover is actually performed lead to the more commonly used composite terms, such as 1+1 and 1:1 protection, which were described in [Section 2.1](#).

i) 1-to-1 Protection

In 1-to-1 protection the working path has a designated recovery path that is only to be used to recover that specific working path.

ii) n-to-1 Protection

In n-to-1 protection, up to n working paths are protected using only one recovery path. If the intent is to protect against any single fault on any of the working paths, the n working paths should be diversely routed between the same PSL and PML. In some cases,

handshaking between PSL and PML may be required to complete the recovery, the details of which are beyond the scope of this draft.

iii) n-to-m Protection

In n-to-m protection, up to n working paths are protected using m recovery paths. Once again, if the intent is to protect against any single fault on any of the n working paths, the n working paths and the m recovery paths should be diversely routed between the same PSL and PML. In some cases, handshaking between PSL and PML may be required to complete the recovery, the details of which are beyond the scope of this draft. -N-to-m protection is for further study.

iv) Split Path Protection

In split path protection, multiple recovery paths are allowed to carry the traffic of a working path based on a certain configurable load splitting ratio. This is especially useful when no single recovery path can be found that can carry the entire traffic of the working path in case of a fault. Split path protection may require handshaking between the PSL and the PML(s), and may require the PML(s) to correlate the traffic arriving on multiple recovery paths with the working path. Although this is an attractive option, the details of split path protection are beyond the scope of this draft, and are for further study.

3.4.3 Bypass Tunnels

It may be convenient, in some cases, to create a "bypass tunnel" for a PPG between a PSL and PML, thereby allowing multiple recovery paths to be transparent to intervening LSRs [8]. In this case, one LSP (the tunnel) is established between the PSL and PML following an acceptable route and a number of recovery paths are supported through the tunnel via label stacking. A bypass tunnel can be used with any of the path mapping options discussed in the previous section.

As with recovery paths, the bypass tunnel may or may not have resource reservations sufficient to provide recovery without service degradation. It is possible that the bypass tunnel may have sufficient resources to recover some number of working paths, but not all at the same time. If the number of recovery paths carrying traffic in the tunnel at any given time is restricted, this is similar to the 1 to n or m to n protection cases mentioned in [Section 3.4.2](#).

3.4.4 Recovery Granularity

Another dimension of recovery considers the amount of traffic requiring protection. This may range from a fraction of a path to a bundle of paths.

3.4.4.1 Selective Traffic Recovery

This option allows for the protection of a fraction of traffic within the same path. The portion of the traffic on an individual path that requires protection is called a protected traffic portion (PTP). A single path may carry different classes of traffic, with different protection requirements. The protected portion of this traffic may be identified by its class, as for example, via the EXP bits in the MPLS shim header or via the priority bit in the ATM header.

3.4.4.2 Bundling

Bundling is a technique used to group multiple working paths together in order to recover them simultaneously. The logical bundling of multiple working paths requiring protection, each of which is routed identically between a PSL and a PML, is called a protected path group (PPG). When a fault occurs on the working path carrying the PPG, the PPG as a whole can be protected either by being switched to a bypass tunnel or by being switched to a recovery path.

3.4.5 Recovery Path Resource Use

In the case of pre-reserved recovery paths, there is the question of what use these resources may be put to when the recovery path is not in use. There are two options:

Dedicated-resource:

If the recovery path resources are dedicated, they may not be used for anything except carrying the working traffic. For example, in the case of 1+1 protection, the working traffic is always carried on the recovery path. Even if the recovery path is not always carrying the working traffic, it may not be possible or desirable to allow other traffic to use these resources.

Extra-traffic-allowed:

If the recovery path only carries the working traffic when the working path fails, then it is possible to allow extra traffic to use the reserved resources at other times. Extra traffic is, by

definition, traffic that can be displaced (without violating service agreements) whenever the recovery path resources are needed for carrying the working path traffic.

3.5 Fault Detection

MPLS recovery is initiated after the detection of either a lower layer fault or a fault at the IP layer or in the operation of MPLS-based mechanisms. We consider four classes of impairments: Path Failure, Path Degraded, Link Failure, and Link Degraded.

Path Failure (PF) is a fault that indicates to an MPLS-based recovery scheme that the connectivity of the path is lost. This may be detected by a path continuity test between the PSL and PML. Some, and perhaps the most common, path failures may be detected using a link probing mechanism between neighbor LSRs. An example of a probing mechanism is a liveness message that is exchanged periodically along the working path between peer LSRs. For either a link probing mechanism or path continuity test to be effective, the test message must be guaranteed to follow the same route as the working or recovery path, over the segment being tested. In addition, the path continuity test must take the path merge points into consideration. In the case of a bi-directional link implemented as two unidirectional links, path failure could mean that either one or both unidirectional links are damaged.

Path Degraded (PD) is a fault that indicates to MPLS-based recovery schemes/mechanisms that the path has connectivity, but that the quality of the connection is unacceptable. This may be detected by a path performance monitoring mechanism, or some other mechanism for determining the error rate on the path or some portion of the path. This is local to the LSR and consists of excessive discarding of packets at an interface, either due to label mismatch or due to TTL errors, for example.

Link Failure (LF) is an indication from a lower layer that the link over which the path is carried has failed. If the lower layer supports detection and reporting of this fault (that is, any fault that indicates link failure e.g., SONET LOS), this may be used by the MPLS recovery mechanism. In some cases, using LF indications may provide faster fault detection than using only MPLS-based fault detection mechanisms.

Link Degraded (LD) is an indication from a lower layer that the link over which the path is carried is performing below an acceptable level. If the lower layer supports detection and reporting of this fault, it may be used by the MPLS recovery mechanism. In some cases,

using LD indications may provide faster fault detection than using only MPLS-based fault detection mechanisms.

3.6 Fault Notification

Protection switching relies on rapid notification of faults. Once a fault is detected, the node that detected the fault must determine if the fault is severe enough to require path recovery. Then the node should send out a notification of the fault by transmitting a FIS to those of its upstream LSRs that were sending traffic on the working path that is affected by the fault. This notification is relayed hop-by-hop by each subsequent LSR to its upstream neighbor, until it eventually reaches a PSL. A PSL is the only LSR that can terminate the FIS and initiate a protection switch of the working path to a recovery path. Since the FIS is a control message, it should be transmitted with high priority to ensure that it propagates rapidly towards the affected PSL(s). Depending on how fault notification is configured in the LSRs of an MPLS domain, the FIS could be sent either as a Layer 2 or Layer 3 packet. An example of a FIS could be the liveness message sent by a downstream LSR to its upstream neighbor, with an optional fault notification field set. Alternatively, it could be a separate fault notification packet. The intermediate LSR should identify which of its incoming links (upstream LSRs) to propagate the FIS on. In the case of 1+1 protection, the FIS should also be sent downstream to the PML where the recovery action is taken.

3.7 Switch-Over Operation

3.7.1 Recovery Trigger

The activation of an MPLS protection switch following the detection or notification of a fault requires a trigger mechanism at the PSL. MPLS protection switching may be initiated due to automatic inputs or external commands. The automatic activation of an MPLS protection switch results from a response to a defect or fault conditions detected at the PSL or to fault notifications received at the PSL. It is possible that the fault detection and trigger mechanisms may be combined, as is the case when a PF, PD, LF, or LD is detected at a PSL and triggers a protection switch to the recovery path. In most cases, however, the detection and trigger mechanisms are distinct, involving the detection of fault at some intermediate LSR followed by the propagation of a fault notification back to the PSL via the FIS, which serves as the protection switch trigger at the PSL. MPLS protection switching in response to external commands results when the operator initiates a protection switch by a command to a PSL (or alternatively by a configuration command to an intermediate LSR, which transmits the FIS towards the PSL).

Note that the PF fault applies to hard failures (fiber cuts, transmitter failures, or LSR fabric failures), as does the LF fault, with the difference that the LF is a lower layer impairment that may be communicated to - MPLS-based recovery mechanisms. The PD (or LD) fault, on the other hand, applies to soft defects (excessive errors due to noise on the link, for instance). The PD (or LD) results in a fault declaration only when the percentage of lost packets exceeds a given threshold, which is provisioned and may be set based on the service level agreement(s) in effect between a service provider and a customer.

3.7.2 Recovery Action

After a fault is detected or FIS is received by the PSL, the recovery action involves either a rerouting or protection switching operation. In both scenarios, the next hop label forwarding entry for a recovery path is bound to the working path.

3.8 Switch-Back Operation

3.8.1 Revertive and Non-Revertive Modes

These protection modes indicate whether or not there is a preferred path for the protected traffic.

3.8.1.1 Revertive Mode

If the working path always is the preferred path, this path will be used whenever it is available. If the working path has a fault, traffic is switched to the recovery path. In the revertive mode of operation, when the preferred path is restored the traffic is automatically switched back to it.

3.8.1.2 Non-revertive Mode

In the non-revertive mode of operation, there is no preferred path. A switchback to the "original" working path is not desired or not possible since the original path may no longer exist after the occurrence of a fault on that path.

If there is a fault on the working path, traffic is switched to the recovery path. When or if the faulty path (the originally working path) is restored, it may become the recovery path (either by configuration, or, if desired, by management actions). This applies for explicitly routed working paths.

When the traffic is switched over to a recovery path, the association between the original working path and the recovery path

may no longer exist, since the original path itself may no longer exist after the fault. Instead, when the network reaches a stable state following routing convergence, the recovery path may be switched over to a different preferred path based either on pre-configured information or optimization based on the new network topology and associated information.

3.8.2 Restoration and Notification

MPLS restoration deals with returning the working traffic from the recovery path to the original or a new working path. Reversion is performed by the PSL upon receiving notification, via FRS, that the working path is repaired or upon receiving notification that a new working path is established.

As before, an LSR that detected the fault on the working path also detects the restoration of the working path. If the working path had experienced a LF defect, the LSR detects a return to normal operation via the receipt of a liveness message from its peer. If the working path had experienced a LD defect at an LSR interface, the LSR could detect a return to normal operation via the resumption of error-free packet reception on that interface. Alternatively, a lower layer that no longer detects a LF defect may inform the MPLS-based recovery mechanisms at the LSR that the link to its peer LSR is operational. The LSR then transmits FRS to its upstream LSR(s) that were transmitting traffic on the working path. This is relayed hop-by-hop until it reaches the PSL(s), at which point the PSL switches the working traffic back to the original working path.

In the non-revertive mode of operation, the working traffic may or may not be restored to the original working path. This is because it might be useful, in some cases, to either: (a) administratively perform a protection switch back to the original working path after gaining further assurances about the integrity of the path, or (b) it may be acceptable to continue operation without the recovery path being protected, or (c) it may be desirable to move the traffic to a new working path that is calculated based on network topology and network policies, after the dynamic routing protocols have converged.

We note that if there is a way to transmit fault information back along a recovery path towards a PSL and if the recovery path is an equivalent recovery path, it is possible for the working path and its recovery path to exchange roles once the original working path is repaired following a fault. This is because, in that case, the recovery path effectively becomes the working path, and the restored working path functions as a recovery path for the original recovery path. This is important, since it affords the benefits of non-

revertive switch operation outlined in [Section 3.8.1](#), without leaving the recovery path unprotected.

[3.8.3](#) Reverting to Preferred Path (or Controlled Rearrangement)

In the revertive mode, a "make before break" restoration switching can be used, which is less disruptive than performing protection switching upon the occurrence of network impairments. This will minimize both packet loss and packet reordering. The controlled rearrangement of paths can also be used to satisfy traffic engineering requirements for load balancing across an MPLS domain.

[3.9](#) Performance

Resource/performance requirements for recovery paths should be specified in terms of the following attributes:

I. Resource class attribute:

Equivalent Recovery Class: The recovery path has the same resource reservations and performance guarantees as the working path. In other words, the recovery path meets the same SLAs as the working path.

Limited Recovery Class: The recovery path does not have the same resource reservations and performance guarantees as the working path.

A. Lower Class: The recovery path has lower resource requirements or less stringent performance requirements than the working path.

B. Best Effort Class: The recovery path is best effort.

II. Priority Attribute:

The recovery path has a priority attribute just like the working path (i.e., the priority attribute of the associated traffic trunks). It can have the same priority as the working path or lower priority.

III. Preemption Attribute:

The recovery path can have the same preemption attribute as the working path or a lower one.

[4.0](#) MPLS Recovery Requirement

The following are the MPLS recovery requirements:

Makam, et al.

I. MPLS recovery SHALL provide an option to identify protection groups (PPGs) and protection portions (PTPs).

II. Each PSL SHALL be capable of performing MPLS recovery upon the detection of the impairments or upon receipt of notifications of impairments.

III. A MPLS recovery method SHALL not preclude manual protection switching commands. This implies that it would be possible under administrative commands to transfer traffic from a working path to a recovery path, or to transfer traffic from a recovery path to a working path, once the working path becomes operational following a fault.

IV. A PSL SHALL be capable of performing either a switch back to the original working path after the fault is corrected or a switchover to a new working path, upon the discovery of a more optimal working path.

V. The recovery model should take into consideration path merging at intermediate LSRs. If a fault affects the merged segment, all the paths sharing that merged segment should be able to recover. Similarly, if a fault affects a non-merged segment, only the path that is affected by the fault should be recovered.

5.0 MPLS Recovery Options

There SHOULD be an option for:

I. Configuration of the recovery path as excess or reserved, with excess as the default. The recovery path that is configured as excess SHALL provide lower priority preemptable traffic access to the protection bandwidth, while the recovery path configured as reserved SHALL not provide any other traffic access to the protection bandwidth.

II. Each protected path SHALL provide an option for configuring the protection alternatives as either rerouting or protection switching.

III. Each protected path SHALL provide a configuration option for enabling restoration as either non-revertive or revertive, with revertive as the default.

IV. Each LSR supporting protection switching SHALL provide an option for fault notification to the PSL.

6.0 Comparison Criteria

Possible criteria to use for comparison of MPLS-based recovery schemes are as follows:

Recovery Time

We define recovery time as the time required for a recovery path to be activated (and traffic flowing) after a fault. Recovery Time is the sum of the Fault Detection Time, Hold-off Time, Notification Time, Recovery Operation Time, and the Traffic Restoration Time. In other words, it is the time between a failure of a node or link in the network and the time before a recovery path is installed and the traffic starts flowing on it.

Full Restoration Time

We define full restoration time as the time required for a permanent restoration. This is the time required for traffic to be routed onto links which are capable of or have been engineered sufficiently to handle traffic in recovery scenarios. Note that this time may or may not be different from the "Recovery Time" depending on whether equivalent or limited recovery paths are used.

Backup Capacity

Recovery schemes may require differing amounts of "backup capacity" in the event of a fault. This capacity will be dependent on the traffic characteristics of the network. However, it may also be dependent on the particular protection plan selection algorithms as well as the signaling and re-routing methods.

Additive Latency

Recovery schemes may introduce additive latency to traffic. For example, a recovery path may take many more hops than the working path. This may be dependent on the recovery path selection algorithms.

Re-ordering

Recovery schemes may introduce re-ordering of packets. Also the action of putting traffic back on preferred paths might cause packet re-ordering.

State Overhead

As the number of recovery paths in a protection plan grows, the state required to maintain them also grows. Schemes may require differing numbers of paths to maintain certain levels of coverage, etc. The state required may also depend on the particular scheme used to recover. In many cases the state overhead will be in proportion to the number of recovery paths.

Loss

Recovery schemes may introduce a certain amount of packet loss during switchover to a recovery path. Schemes that introduce loss during recovery can measure this loss by evaluating recovery times in proportion to the link speed.

In case of link or node failure a certain packet loss is inevitable.

Coverage

Recovery schemes may offer various types of failover coverage. The total coverage may be defined in terms of several metrics:

I. Fault Types: Recovery schemes may account for only link faults or both node and link faults or also degraded service. For example, a scheme may require more recovery paths to take node faults into account.

II. Number of concurrent faults: dependent on the layout of recovery paths in the protection plan, multiple fault scenarios may be able to be restored.

III. Number of recovery paths: for a given fault, there may be one or more recovery paths.

IV. Percentage of coverage: dependent on a scheme and its implementation, a certain percentage of faults may be covered. This may be subdivided into percentage of link faults and percentage of node faults.

V. The number of protected paths may effect how fast the total set of paths affected by a fault could be recovered. The ratio of protected is n/N , where n is the number of protected paths and N is the total number of paths.

7.0 Security Considerations

The MPLS recovery that is specified herein does not raise any security issues that are not already present in the MPLS architecture.

8.0 Intellectual Property Considerations

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

9.0 Acknowledgements

We would like to thank members of the MPLS WG mailing list for their suggestions on the earlier version of this draft. In particular, Bora Akyol, Dave Allan, and Neil Harrisson, whose suggestions and comments were very helpful in revising the document.

10.0 Authors' Addresses

Srinivas Makam
Tellabs Operations, Inc.
4951 Indiana Avenue
Lisle, IL 60532
Phone: 630-512-7217
Srinivas.Makam@tellabs.com

Vishal Sharma
Tellabs Research Center
One Kendall Square
Bldg. 100, Ste. 121
Cambridge, MA 02139-1562
Phone: 617-577-8760
Vishal.Sharma@tellabs.com

Ken Owens
Tellabs Operations, Inc.
1106 Fourth Street
St. Louis, MO 63126
Phone: 314-918-1579
Ken.Owens@tellabs.com

Changcheng Huang
Tellabs Operations, Inc.
4951 Indiana Avenue
Lisle, IL 60532
Phone: 630-512-7754
Changcheng.Huang@tellabs.com

Ben Mack-Crane
Tellabs Operations, Inc.
4951 Indiana Avenue
Lisle, IL 60532
Ph: 630-512-7255
Ben.Mack-Crane@tellabs.com

Fiffi Hellstrand
Nortel Networks
St Eriksgatan 115, PO Box 6701
113 85 Stockholm, Sweden
Ph: +46 8 5088 3687
Fiffi@nortelnetworks.com

Jon Weil
Nortel Networks
Harlow Laboratories London Road
Harlow Essex CM17 9NA, UK
Phone: +44 (0)1279 403935
jonweil@nortelnetworks.com

Brad Cain
Mirror Image Internet
49 Dragon Ct.
Woburn, MA 01801, USA
bcain@mirror-image.com

Loa Andersson
Nortel Networks
St Eriksgatan 115, PO Box 6701
113 85 Stockholm, Sweden
phone: +46 8 50 88 36 34
loa.andersson@nortelnetworks.com

Bilel Jamoussi
Nortel Networks
3 Federal Street, BL3-03
Billerica, MA 01821, USA
jamoussi@nortelnetworks.com

Seyhan Civanlar
Coreon, Inc.
1200 South Avenue, Suite 103
Staten Island, NY 10314
Ph: (718) 889 4203
scivanlar@coreon.net

Angela Chiu
AT&T Labs, Rm. 4-204,
100 Schulz Dr.
Red Bank, NJ 07701
Ph: (732) 345-3441
alchiu@att.com

11.0 References

- [1] Rosen, E., Viswanathan, A., and Callon, R., "Multiprotocol Label Switching Architecture", Work in Progress, Internet Draft <[draft-ietf-mpls-arch-06.txt](#)>, August 1999.
- [2] Andersson, L., Doolan, P., Feldman, N., Fredette, A., Thomas, B., "LDP Specification", Work in Progress, Internet Draft <[draft-ietf-mpls-ldp-06.txt](#)>, September 1999.
- [3] Awduche, D. Hannan, A., and Xiao, X., "Applicability Statement for Extensions to RSVP for LSP-Tunnels", [draft-ietf-mpls-rsvp-tunnel-applicability-00.txt](#), work in progress, Sept. 1999.
- [4] Jamoussi, B. "Constraint-Based LSP Setup using LDP", Work in Progress, Internet Draft <[draft-ietf-mpls-cr-ldp-03.txt](#)>, September 1999.
- [5] Braden, R., Zhang, L., Berson, S., Herzog, S., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", [RFC 2205](#), September 1997.
- [6] Awduche, D. et al "Extensions to RSVP for LSP Tunnels", Work in Progress, Internet Draft <[draft-ietf-mpls-rsvp-lsp-tunnel-04.txt](#)>, September 1999.
- [7] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., McManus, J., "Requirements for Traffic Engineering Over MPLS", [RFC 2702](#), September 1999.
- [8] Andersson, L., Cain B., Jamoussi, B., "Requirement Framework for Fast Re-route with MPLS", [draft-andersson-reroute-frmwrk-00.txt](#), work in progress, October 1999.
- [9] Goguen, R. and Swallow, G., "RSVP Label Allocation for Backup Tunnels", [draft-swallow-rsvp-bypass-label-00.txt](#), work in progress, October 1999.
- [10] Makam, S., Sharma, V., Owens, K., Huang, C., "Protection/restoration of MPLS Networks", [draft-makam-mpls-protection-00.txt](#), work in progress, October 1999.
- [11] Callon, R., Doolan, P., Feldman, N., Fredette, A., Swallow, G., Viswanathan, A., "A Framework for Multiprotocol Label Switching", <[draft-ietf-mpls-framework-05.txt](#)>, Work in Progress, September 1999.

- [12] Haskin, D. and Krishnan R., "A Method for Setting an Alternative Label Switched Path to Handle Fast Reroute", [draft-haskin-mpls-fast-reroute-01.txt](#), 1999, Work in progress.

