             Recovery (Protection and Restoration) Terminology
            for Generalized Multi-Protocol Label Switching (GMPLS)

               draft-ietf-ccamp-gmpls-recovery-terminology-06.txt

Status of this Memo

Copyright Notice

Abstract

    This document defines a common terminology for Generalized Multi-
    Protocol Label Switching (GMPLS) based recovery mechanisms (i.e.
    protection and restoration). The terminology is independent of the
    underlying transport technologies covered by GMPLS.

Table of Contents

1. Contributors

   This document is the result of the CCAMP Working Group Protection
   and Restoration design team joint effort. The following are the
   authors that contributed to the present document:

   Deborah Brungard (AT&T)
   Rm. D1-3C22 - 200 S. Laurel Ave.
   Middletown, NJ 07748, USA
   EMail: dbrungard@att.com

   Sudheer Dharanikota
   EMail: sudheer@ieee.org

   Jonathan P. Lang (Sonos)
   506 Chapala Street
   Santa Barbara, CA 93101, USA
   EMail: jplang@ieee.org

   Guangzhi Li (AT&T)
   180 Park Avenue,
   Florham Park, NJ 07932, USA
   EMail: gli@research.att.com

   Eric Mannie
   EMail: eric_mannie@hotmail.com

   Dimitri Papadimitriou (Alcatel)
   Francis Wellesplein, 1
   B-2018 Antwerpen, Belgium
   EMail: dimitri.papadimitriou@alcatel.be

   Bala Rajagopalan (Intel Broadband Wireless Division)

2111 NE 25th Ave.
Hillsboro, OR 97124, USA
EMail: bala.rajagopalan@intel.com

Yakov Rekhter (Juniper)
1194 N. Mathilda Avenue
Sunnyvale, CA 94089, USA
EMail: yakov@juniper.net

2. Introduction

This document defines a common terminology for Generalized Multi-Protocol Label Switching (GMPLS) based recovery mechanisms (i.e. protection and restoration).

The terminology proposed in this document is independent of the underlying transport technologies and borrows from the G.808.1 ITU-T Recommendation [G.808.1] and from the G.841 ITU-T Recommendation [G.841]. The restoration terminology and concepts have been gathered from numerous sources including IETF drafts.

In the context of this document, the term "recovery" denotes both protection and restoration. The specific terms "protection" and "restoration" will only be used when differentiation is required.

This document focuses on the terminology for the recovery of Label Switched Paths (LSPs) controlled by a GMPLS control plane. The proposed terminology applies to end-to-end, segment, and span (i.e. link) recovery. Note that the terminology for recovery of the control plane itself is not in the scope of this document.

Protection and restoration of switched LSPs under tight time constraints is a challenging problem. This is particularly relevant to optical networks that consist of Time Division Multiplex (TDM) and/or all-optical (photonic) cross-connects referred to as GMPLS nodes (or simply nodes, or even sometimes "Label Switching Routers, or LSRs") connected in a general topology [RFC3945].

Recovery typically involves the activation of a recovery (or alternate) LSP when a failure is encountered in the working LSP.

A working or recovery LSP is characterized by an ingress interface, an egress interface, and a set of intermediate nodes and spans through which the LSP is routed. The working and recovery LSPs are

typically resource disjoint (e.g. node and/or span disjoint). This
ensures that a single failure will not affect both the working and
recovery LSPs.

A bi-directional span between neighboring nodes is usually realized
as a pair of unidirectional spans. The end-to-end path for a bi-
directional LSP therefore consists of a series of bi-directional
segments (i.e. Sub-Network Connections, or SNCs, in the ITU-T
terminology) between the source and destination nodes, traversing
intermediate nodes.

3. Conventions used in this document:

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in
this document are to be interpreted as described in [RFC2119].

4. Recovery Terminology Common to Protection and Restoration

This section defines the following general terms common to both
protection and restoration (i.e. recovery). In addition, most of
these terms apply to end-to-end, segment and span LSP recovery. Note

that span recovery does not protect the nodes at each end of the
span, otherwise end-to-end or segment LSP recovery should be used.

The terminology and the definitions have been originally taken from
[G.808.1]. However, for generalization, the following language that
is not directly related to recovery has been adapted to GMPLS and
the common IETF terminology:

An LSP is used as a generic term to designate either an SNC (Sub-
Network Connection) or an NC (Network Connection) in ITU-T
terminology. The ITU-T uses the term transport entity to designate
either a link, an SNC or an NC. The term "Traffic" is used instead
of "Traffic Signal". The term protection or restoration "scheme" is
used instead of protection or restoration "architecture".

The reader is invited to read [G.841] and [G.808.1] for references
to SDH protection and Generic Protection Switching terminology,
respectively. Note that restoration is not in the scope of
[G.808.1].

4.1 Working and Recovery LSP/Span

A working LSP/span is an LSP/span transporting "normal" user
traffic. A recovery LSP/span is an LSP/span used to transport
"normal" user traffic when the working LSP/span fails. Additionally,
the recovery LSP/span may transport "extra" user traffic (i.e. pre-
emptable traffic) when normal traffic is carried over the working
LSP/span.

4.2 Traffic Types

The different types of traffic that can be transported over an
LSP/span in the context of this document are defined hereafter:

A. Normal traffic:

User traffic that may be protected by two alternative LSPs/spans
(the working and recovery LSPs/spans).

B. Extra traffic:

User traffic carried over recovery resources (e.g. a recovery
LSP/span) when these resources are not being used for the recovery
of normal traffic; i.e. when the recovery resources are in standby
mode. When the recovery resources are required to recover normal
traffic from the failed working LSP/span, the extra traffic is pre-
empted. Extra traffic is not protected by definition, but may be
restored. Moreover, extra traffic does not need to commence or be
terminated at the ends of the LSPs/spans that it uses.

C. Null traffic:

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

Traffic carried over the recovery LSP/span if it is not used to
carry normal or extra traffic. Null traffic can be any kind of
traffic that conforms to the signal structure of the specific layer,
and it is ignored (not selected) at the egress of the recovery
LSP/span.

4.3 LSP/Span Protection and Restoration

The following subtle distinction is generally made between the terms
"protection" and "restoration", even though these terms are often
used interchangeably [RFC3386].

The distinction between protection and restoration is made based on
the resource allocation done during the recovery LSP/span
establishment. The distinction between different types of

restoration is made based on the level of route computation,
signaling and resource allocation done during the restoration
LSP/span establishment.

A. LSP/Span Protection

LSP/span protection denotes the paradigm whereby one or more
dedicated protection LSP(s)/span(s) is/are fully established to
protect one or more working LSP(s)/span(s).

For a protection LSP, this implies that route computation took
place, that the LSP was fully signaled all the way and that its
resources were fully selected (i.e. allocated) and cross-connected
between the ingress and egress nodes.

For a protection span, this implies that the span has been selected
and reserved for protection.

Indeed, it means that no signaling takes place to establish the
protection LSP/span when a failure occurs. However, various other
kinds of signaling may take place between the ingress and egress
nodes for fault notification, to synchronize their use of the
protection LSP/span, for reversion, etc.

B. LSP/Span Restoration

LSP/span restoration denotes the paradigm whereby some restoration
resources may be pre-computed, signaled and selected a priori, but
not cross-connected to restore a working LSP/span. The complete
establishment of the restoration LSP/span occurs only after a
failure of the working LSP/span, and requires some additional
signaling.

Both protection and restoration require signaling. Signaling to
establish the recovery resources and signaling associated with the
use of the recovery LSP(s)/span(s) are needed.

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

4.4 Recovery Scope

Recovery can be applied at various levels throughout the network. An
LSP may be subject to local (span), segment, and/or end-to-end
recovery.

Local (span) recovery refers to the recovery of an LSP over a link
between two nodes.

End-to-end recovery refers to the recovery of an entire LSP from its
source (ingress node end-point) to its destination (egress node end-
point).

Segment recovery refers to the recovery over a portion of the
network of a segment LSP (i.e. an SNC in the ITU-T terminology) of
an end-to-end LSP. Such recovery protects against span and/or node
failure over a particular portion of the network traversed by an
end-to-end LSP.

4.5 Recovery Domain

A recovery domain is defined as a set of nodes and spans over which
one or more recovery schemes are provided. A recovery domain served
by one single recovery scheme is referred to as a "single recovery
domain", while a recovery domain served by multiple recovery schemes
is referred to as a "multi recovery domain".

The recovery operation is contained within the recovery domain. A
GMPLS recovery domain must be entirely contained within a GMPLS
domain. A GMPLS domain (defined as a set of nodes and spans
controlled by GMPLS) may contain multiple recovery domains.

4.6 Recovery Types

The different recovery types can be classified depending on the
number of recovery LSPs/spans that are protecting a given number of
working LSPs/spans. The definitions given hereafter are from the
point of view of a working LSP/span that needs to be protected by a
recovery scheme.

A. 1+1 type: dedicated protection

One dedicated protection LSP/span protects exactly one working
LSP/span and the normal traffic is permanently duplicated at the
ingress node on both the working and protection LSPs/spans. No extra
traffic can be carried over the protection LSP/span.

This type is applicable to LSP/span protection, but not to LSP/span
restoration.

B. 0:1 type: unprotected

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

No specific recovery LSP/span protects the working LSP/span.

However, the working LSP/span can potentially be restored through any alternate available route/span, with or without any pre-computed restoration route. Note that there are no resources pre-established for this recovery type.

This type is applicable to LSP/span restoration, but not to LSP/span protection. Span restoration can be for instance achieved by moving all the LSPs transported over of a failed span to a dynamically selected span.

C. 1:1 type: dedicated recovery with extra traffic

One specific recovery LSP/span protects exactly one specific working LSP/span but the normal traffic is transmitted only over one LSP (working or recovery) at a time. Extra traffic can be transported using the recovery LSP/span resources.

This type is applicable to LSP/span protection and LSP restoration, but not to span restoration.

D. 1:N (N > 1) type: shared recovery with extra traffic

A specific recovery LSP/span is dedicated to the protection of up to N working LSPs/spans. The set of working LSPs/spans is explicitly identified. Extra traffic can be transported over the recovery LSP/span. All these LSPs/spans must start and end at the same nodes.

Sometimes, the working LSPs/spans are assumed to be resource disjoint in the network so that they do not share any failure probability, but this is not mandatory. Obviously, if more than one working LSP/span in the set of N are affected by some failure(s) at the same time, the traffic on only one of these failed LSPs/spans may be recovered over the recovery LSP/span. Note that N can be arbitrarily large (i.e. infinite). The choice of N is a policy decision.

This type is applicable to LSP/span protection and LSP restoration, but not to span restoration.

Note: a shared recovery where each recovery resource can be shared by a maximum of X LSPs/spans is not defined as a recovery type but as a recovery scheme. The choice of X is a network resource management policy decision.

E. M:N (M, N > 1, N >= M) type:

A set of M specific recovery LSPs/spans protects a set of up to N specific working LSPs/spans. The two sets are explicitly identified. Extra traffic can be transported over the M recovery LSPs/spans when available. All the LSPs/spans must start and end at the same nodes.

   Sometimes, the working LSPs/spans are assumed to be resource
   disjoint in the network so that they do not share any failure
   probability, but this is not mandatory. Obviously, if several
   working LSPs/spans in the set of N are concurrently affected by some
   failure(s), the traffic on only M of these failed LSPs/spans may be
   recovered. Note that N can be arbitrarily large (i.e. infinite). The
   choice of N and M is a policy decision.

   This type is applicable to LSP/span protection and LSP restoration,
   but not to span restoration.

4.7 Bridge Types

   A bridge is the function that connects the normal traffic and extra
   traffic to the working and recovery LSP/span.

   A. Permanent bridge

   Under a 1+1 type, the bridge connects the normal traffic to both the
   working and protection LSPs/spans. This type of bridge is not
   applicable to restoration types. There is of course no extra traffic
   connected to the recovery LSP/span.

   B. Broadcast bridge

   For 1:N and M:N types, the bridge permanently connects the normal
   traffic to the working LSP/span. In the event of recovery switching,
   the normal traffic is additionally connected to the recovery
   LSP/span. Extra traffic is either not connected or connected to the
   recovery LSP/span.

   C. Selector bridge

   For 1:N and M:N types, the bridge connects the normal traffic to
   either the working or the recovery LSP/span. Extra traffic is either
   not connected or connected to the recovery LSP/span.

4.8 Selector Types

   A selector is the function that extracts the normal traffic either
   from the working or the recovery LSP/span. Extra traffic is either
   extracted from the recovery LSP/span, or is not extracted.

   A. Selective selector

   Is a selector that extracts the normal traffic from either the

working LSP/span output or the recovery LSP/span output.

B. Merging selector

For 1:N and M:N protection types, the selector permanently extracts
the normal traffic from both the working and recovery LSP/span

outputs. This alternative works only in combination with a selector
bridge.

4.9 Recovery GMPLS Nodes

This section defines the GMPLS nodes involved during recovery.

A. Ingress GMPLS node of an end-to-end LSP/segment LSP/span

The ingress node of an end-to-end LSP/segment LSP/span is where the
normal traffic may be bridged to the recovery end-to-end LSP/segment
LSP/span. Also known as source node in the ITU-T terminology.

B. Egress GMPLS node of an end-to-end LSP/segment LSP/span

The egress node of an end-to-end LSP/segment LSP/span is where the
normal traffic may be selected from either the working or the
recovery end-to-end LSP/segment LSP/span. Also known as sink node in
the ITU-T terminology.

C. Intermediate GMPLS node of an end-to-end LSP/segment LSP

A node along either the working or recovery end-to-end LSP/segment
LSP route between the corresponding ingress and egress nodes. Also
known as intermediate node in the ITU-T terminology.

4.10 Switch-over Mechanism

A switch-over is an action that can be performed at both the bridge
and the selector. This action is as follows:

A. For the selector:

The action of selecting normal traffic from the recovery LSP/span
rather than from the working LSP/span.

B. For the bridge:

In case of permanent connection to the working LSP/span, the action
of connecting or disconnecting the normal traffic to the recovery

LSP/span. In case of non-permanent connection to the working
LSP/span, the action of connecting the normal traffic to the
recovery LSP/span.

4.11 Reversion operations

A revertive recovery operation refers to a recovery switching
operation, where the traffic returns to (or remains on) the working
LSP/span if the switch-over requests are terminated; i.e. when the
working LSP/span has recovered from the failure.

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

Therefore a non-revertive recovery switching operation is when the
traffic does not return to the working LSP/span if the switch-over
requests are terminated.

4.12 Failure Reporting

This section gives (for information) several signal types commonly
used in transport planes to report a failure condition. Note that
fault reporting may require additional signaling mechanisms.

A. Signal Degrade (SD): a signal indicating that the associated data
has degraded.

B. Signal Fail (SF): a signal indicating that the associated data
has failed.

C. Signal Degrade Group (SDG): a signal indicating that the
associated group data has degraded.

D. Signal Fail Group (SFG): a signal indicating that the associated
group has failed.

Note: SDG and SFG definitions are under discussion at the ITU-T.

4.13 External commands

This section defines several external commands, typically issued by
an operator through the Network Management System (NMS)/Element
Management System (EMS), which can be used to influence or command
the recovery schemes.

A. Lockout of recovery LSP/span:

A configuration action initiated externally that results in the
recovery LSP/span being temporarily unavailable to transport traffic
(either normal or extra traffic).

B. Lockout of normal traffic:

A configuration action initiated externally that results in the
normal traffic being temporarily not allowed to be routed over its
recovery LSP/span. Note that in this case extra-traffic is still
allowed on the recovery LSP/span.

C. Freeze:

A configuration action initiated externally that prevents any
switch-over action to be taken, and as such freezes the current
state.

D. Forced switch-over for normal traffic:

A switch-over action initiated externally that switches normal
traffic to the recovery LSP/span, unless an equal or higher priority
switch-over command is in effect.

E. Manual switch-over for normal traffic:

A switch-over action initiated externally that switches normal
traffic to the recovery LSP/span, unless a fault condition exists on
other LSPs/spans (including the recovery LSP/span) or an equal or
higher priority switch-over command is in effect.

F. Manual switch-over for recovery LSP/span:

A switch-over action initiated externally that switches normal
traffic to the working LSP/span, unless a fault condition exists on
the working LSP/span or an equal or higher priority switch-over
command is in effect.

G. Clear:

An action initiated externally that clears the active external
command.

4.14 Unidirectional versus Bi-Directional Recovery Switching

A. Unidirectional recovery switching:

A recovery switching mode in which, for a unidirectional fault (i.e.
a fault affecting only one direction of transmission), only the
normal traffic transported in the affected direction (of the LSP or
span) is switched to the recovery LSP/span.

B. Bi-directional recovery switching:

A recovery switching mode in which, for a unidirectional fault, the
normal traffic in both directions (of the LSP or span), including
the affected direction and the unaffected direction, are switched to
the recovery LSP/span.

4.15 Full versus Partial Span Recovery Switching

Bulk LSP recovery is initiated upon reception on either span failure
notification or bulk failure notification of the S LSPs carried by
this span. In either case, the corresponding recovery switching
actions are performed at the LSP level such that the ratio between
the number of recovery switching messages and the number of
recovered LSP (in one given direction) is minimized. If this ratio
equals 1, one refers to full span recovery, otherwise, if this ratio
is greater than 1 one refers to partial span recovery.

A. Full Span Recovery

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

All the S LSP carried over a given span are recovered under span
failure condition. Full span recovery is also referred to as "bulk
recovery".

B. Partial Span Recovery

Only a subset s of the S LSP carried over a given span are recovered
under span failure condition. Both selection criteria of the
entities belonging to this subset and the decision concerning the
recovery of the remaining (S - s) LSP are based on local policy.

4.16 Recovery Schemes Related Time and Durations

This section gives several typical timing definitions that are of
importance for recovery schemes.

A. Detection time:

The time between the occurrence of the fault or degradation and its

detection. Note that this is a rather theoretical time since in
practice this is difficult to measure.

   B. Correlation time:

   The time between the detection of the fault or degradation and the
   reporting of the signal fail or degrade. This time is typically used
   in correlating related failures or degradations.

   C. Notification time:

   The time between the reporting of the signal fail or degrade and the
   reception of the indication of this event by the entities that
   decide on the recovery switching operation(s).

   D. Recovery Switching time:

   The time between the initialization of the recovery switching
   operation and the moment the normal traffic is selected from the
   recovery LSP/span.

   E. Total Recovery time:

   The total recovery time is defined as the sum of the detection, the
   correlation, the notification and the recovery switching time.

   F. Wait To Restore time:

   A period of time that must elapse from a recovered fault before an
   LSP/span can be used again to transport the normal traffic and/or to
   select the normal traffic from.

   Note: the hold-off time is defined as the time between the reporting
   of signal fail or degrade, and the initialization of the recovery
   switching operation. This is useful when multiple layers of recovery
   are being used.

4.17 Impairment

   A defect or performance degradation, which may lead to SF or SD
   trigger.

4.18 Recovery Ratio

   The quotient of the actually recovery bandwidth divided by the

traffic bandwidth which is intended to be protected.

4.19 Hitless Protection Switch-over

Protection switch-over, which does not cause data loss, data
duplication, data disorder, or bit errors upon recovery switching
action.

4.20 Network Survivability

The set of capabilities that allow a network to restore affected
traffic in the event of a failure. The degree of survivability is
determined by the network's capability to survive single and
multiple failures.

4.21 Survivable Network

A network that is capable of restoring traffic in the event of a
failure.

4.22 Escalation

A network survivability action caused by the impossibility of the
survivability function in lower layers.

5. Recovery Phases

It is commonly accepted that recovery implies that the following
generic operations need to be performed when an LSP/span or a node
failure occurs:

- Phase 1: Failure Detection

The action of detecting the impairment (defect of performance
degradation) as a defect condition and consequential activation of
SF or SD trigger to the control plane (through internal interface
with the transport plane). Thus, failure detection (that should
occur at the transport layer closest to the failure) is the only
phase that can not be achieved by the control plane alone.

- Phase 2: Failure Localization (and Isolation)

Failure localization provides to the deciding entity information
about the location (and so the identity) of the transport plane
entity that causes the LSP(s)/span(s) failure. The deciding entity

can then take accurate decision to achieve finer grained recovery
switching action(s).

- Phase 3: Failure Notification

Failure notification phase is used 1) to inform intermediate nodes
that LSP(s)/span(s) failure has occurred and has been detected 2) to
inform the recovery deciding entities (which can correspond to any
intermediate or end-point of the failed LSP/span) that the
corresponding LSP/span is not available.

- Phase 4: Recovery (Protection or Restoration)

See Section 4.3.

- Phase 5: Reversion (Normalization)

See Section 4.11.

The combination of Failure Detection and Failure Localization and
Notification is referred to as Fault Management.

5.1 Entities Involved During Recovery

The entities involved during the recovery operations can be defined
as follows; these entities are parts of ingress, egress and
intermediate nodes as defined previously:

A. Detecting Entity (Failure Detection):

An entity that detects a failure or group of failures; providing
thus a non-correlated list of failures.

B. Reporting Entity (Failure Correlation and Notification):

An entity that can make an intelligent decision on fault correlation
and report the failure to the deciding entity. Fault reporting can
be automatically performed by the deciding entity detecting the
failure.

C. Deciding Entity (part of the failure recovery decision process):

An entity that makes the recovery decision or selects the recovery
resources. This entity communicates the decision to the impacted
LSPs/spans with the recovery actions to be performed.

D. Recovering Entity (part of the failure recovery activation process):

An entity that participates in the recovery of the LSPs/spans.

The process of moving failed LSPs from a failed (working) span to a protection span must be initiated by one of the nodes terminating the span, e.g. A or B. The deciding (and recovering) entity is referred to as the "master" while the other node is called the "slave" and corresponds to a recovering only entity.

Note: The determination of the master and the slave may be based on configured information or protocol specific requirements.

6. Protection Schemes

This section clarifies the multiple possible protection schemes and the specific terminology for the protection.

6.1 1+1 Protection

1+1 protection has one working LSP/span, one protection LSP/span and a permanent bridge. At the ingress node, the normal traffic is permanently bridged to both the working and protection LSP/span. At the egress node, the normal traffic is selected from the better of the two LSPs/spans.

Due to the permanent bridging, the 1+1 protection does not allow an unprotected extra traffic signal to be provided.

6.2 1:N (N >= 1) Protection

1:N protection has N working LSPs/spans carrying normal traffic and 1 protection LSP/span that may carry extra-traffic.

At the ingress, the normal traffic is either permanently connected to its working LSP/span and may be connected to the protection LSP/span (case of broadcast bridge), or is connected to either its working or the protection LSP/span (case of selector bridge). At the egress node, the normal traffic is selected from either its working or protection LSP/span.

Unprotected extra traffic can be transported over the protection LSP/span whenever the protection LSP/span is not used to carry a normal traffic.

6.3 M:N (M, N > 1, N >= M) Protection

M:N protection has N working LSPs/spans carrying normal traffic and M protection LSP/span that may carry extra-traffic.

   At the ingress, the normal traffic is either permanently connected
   to its working LSP/span and may be connected to one of the
   protection LSPs/spans (case of broadcast bridge), or is connected to
   either its working or one of the protection LSPs/spans (case of
   selector bridge). At the egress node, the normal traffic is selected
   from either its working or one of the protection LSP/span.

   Unprotected extra traffic can be transported over the M protection
   LSP/span whenever the protection LSPs/spans is not used to carry a
   normal traffic.

6.4 Notes on Protection Schemes

   All protection types are either uni- or bi-directional, obviously,
   the latter applies only to bi-directional LSP/span and requires
   coordination between the ingress and egress node during protection
   switching.

   All protection types except 1+1 unidirectional protection switching
   require a communication channel between the ingress and the egress
   node.

   In the GMPLS context, span protection refers to the full or partial
   span recovery of the LSPs carried over that span (see Section 4.15).

7. Restoration Schemes

   This section clarifies the multiple possible restoration schemes and
   the specific terminology for the restoration.

7.1 Pre-planned LSP Restoration

   Also referred to as pre-planned LSP re-routing. Before failure
   detection and/or notification, one or more restoration LSPs are
   instantiated between the same ingress-egress node pair than the
   working LSP. Note that the restoration resources must be pre-
   computed, must be signaled and may be selected a priori, but not
   cross-connected. Thus, the restoration LSP is not able to carry any
   extra-traffic.

   The complete establishment of the restoration LSP (i.e. activation)
   occurs only after failure detection and/or notification of the
   working LSP and requires some additional restoration signaling.
   Therefore, this mechanism protects against working LSP failure(s)
   but requires activation of the restoration LSP after failure

occurrence. After the ingress node has activated the restoration LSP, the latter can carry the normal traffic.

Note: when each working LSP is recoverable by exactly one restoration LSP, one refers also to 1:1 (pre-planned) re-routing without extra-traffic.

7.1.1 Shared-Mesh Restoration

"Shared-mesh" restoration is defined as a particular case of pre-planned LSP re-routing that reduces the restoration resource requirements by allowing multiple restoration LSPs (initiated from distinct ingress nodes) to share common resources (including links and nodes.)

7.2 LSP Restoration

Also referred to as LSP re-routing. The ingress node switches the normal traffic to an alternate LSP signaled and fully established (i.e. cross-connected) after failure detection and/or notification. The alternate LSP path may be computed after failure detection and/or notification. In this case, one also refers to "Full LSP Re-routing."

The alternate LSP is signaled from the ingress node and may reuse intermediate node's resources of the working LSP under failure condition (and may also include additional intermediate nodes.)

7.2.1 Hard LSP Restoration

Also referred to as hard LSP re-routing. A re-routing operation where the LSP is released before the full establishment of an alternate LSP (i.e. break-before-make).

7.2.2 Soft LSP Restoration

Also referred to as soft LSP re-routing. A re-routing operation where the LSP is released after the full establishment of an alternate LSP (i.e. make-before-break).

8. Security Considerations

Security considerations are detailed in [ANAL] and [FUNCT].

9. IANA Considerations

This document defines no new code points and requires no action by
IANA.

10. References

10.1 Normative References

[ANAL]   D.Papadimitriou and E.Mannie (Editors), "Analysis of
                Generalized Multi-Protocol Label Switching (GMPLS)-
                based Recovery Mechanisms (including Protection and
                Restoration)", Internet Draft (Work in progress), April
                2005.

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

[FUNCT]   J.P.Lang, B.Rajagopalan and D.Papadimitriou (Editors),
                "Generalized MPLS Recovery Functional Specification,"
                Internet Draft (Work in progress), April 2005.

[RFC2026]   S.Bradner, "The Internet Standards Process -- Revision
                3", BCP 9, RFC 2026, October 1996.

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

[RFC3667]   S.Bradner, "IETF Rights in Contributions", BCP 78,
                RFC 3667, February 2004.

[RFC3668]   S.Bradner, Ed., "Intellectual Property Rights in IETF
                Technology", BCP 79, RFC 3668, February 2004.

10.2 Informative References

[RFC3386]   W.S.Lai, et al., "Network Hierarchy and Multilayer
                Survivability," RFC 3386, November 2002.

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

[T1.105]   ANSI, "Synchronous Optical Network (SONET): Basic
                Description Including Multiplex Structure, Rates, and
                Formats," ANSI T1.105, January 2001.

For information on the availability of the following documents,
please see http://www.itu.int

   [G.707]   ITU-T, "Network Node Interface for the Synchronous
                Digital Hierarchy (SDH)," Recommendation G.707, October
                2000.

   [G.783]   ITU-T, "Characteristics of Synchronous Digital
                Hierarchy (SDH) Equipment Functional Blocks,"
                Recommendation G.783, October 2000.

   [G.806]   ITU-T, "Characteristics of Transport Equipment -
                Description Methodology and Generic Functionality,"
                Recommendation G.806, October 2000.

   [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.

E.Mannie, D.Papadimitriou et al.- Expires October 2005   19

draft-ietf-ccamp-gmpls-recovery-terminology-06.txt   March 2005

   [G.842]   ITU-T, "Interworking of SDH network protection
                architectures," Recommendation G.842, October 1998.

11. **Acknowledgments**

   Many thanks to Adrian Farrel for having thoroughly review this
   document.

12. **Editor's Addresses**

   Eric Mannie
   EMail: eric_mannie@hotmail.com

   Dimitri Papadimitriou
   Alcatel
   Francis Wellesplein, 1
   B-2018 Antwerpen, Belgium
   Phone: +32 3 240-8491
   EMail: dimitri.papadimitriou@alcatel.be

Intellectual Property Statement

   The IETF takes no position regarding the validity or scope of any
   Intellectual Property Rights or other rights that might be claimed
   to pertain to the implementation or use of the technology described
   in this document or the extent to which any license under such
   rights might or might not be available; nor does it represent that
   it has made any independent effort to identify any such rights.
   Information on the procedures with respect to rights in RFC
   documents can be found in BCP 78 and BCP 79.

   Copies of IPR disclosures made to the IETF Secretariat and any
   assurances of licenses to be made available, or the result of an
   attempt made to obtain a general license or permission for the use
   of such proprietary rights by implementers or users of this
   specification can be obtained from the IETF on-line IPR repository
   at http://www.ietf.org/ipr.

   The IETF invites any interested party to bring to its attention any
   copyrights, patents or patent applications, or other proprietary

rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Disclaimer of Validity

Copyright Statement

Acknowledgment