**RSVP-TE Extensions in support of End-to-End
Generalized Multi-Protocol Label Switching (GMPLS) Recovery**

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt

Status of this Memo

Copyright Notice

Abstract

   This document describes protocol specific procedures and extensions
   for Generalized Multi-Protocol Label Switching (GMPLS) Resource
   ReserVation Protocol - Traffic Engineering (RSVP-TE) signaling to
   support end-to-end Label Switched Path (LSP) recovery that denotes
   protection and restoration. A generic functional description of
   GMPLS recovery can be found in a companion document, RFC 4426.

J.P.Lang et al.   Standards Track   1

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Table of Contents

J.P.Lang et al.   Expires February 2007   2

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

1. 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].

   In addition, the reader is assumed to be familiar with the
   terminology used in [RFC3945], [RFC3471], [RFC3473] and referenced
   as well as in [RFC4427] and [RFC4426].

2. Introduction

   Generalized Multi-Protocol Label Switching (GMPLS) extends MPLS to
   include support for Layer-2 (L2SC), Time-Division Multiplex (TDM),
   Lambda Switch Capable (LSC), and Fiber Switch Capable (FSC)
   interfaces. GMPLS recovery uses control plane mechanisms (i.e.,
   signaling, routing, link management mechanisms) to support data
   plane fault recovery. Note that the analogous (data plane) fault
   detection mechanisms are required to be present in support of the
   control plane mechanisms. In this document, the term "recovery" is
   generically used to denote both protection and restoration; the
   specific terms "protection" and "restoration" are only used when
   differentiation is required. The subtle distinction between
   protection and restoration is made based on the resource allocation
   done during the recovery phase (see [RFC4427]).

   A functional description of GMPLS recovery is provided in [RFC4426]
   and should be considered as a companion document. The present
   document describes the protocol specific procedures for GMPLS RSVP-
   TE (Resource ReSerVation Protocol - Traffic Engineering) signaling
   (see [RFC3473]) to support end-to-end recovery. End-to-end recovery

refers to the recovery of an entire LSP from its head-end (ingress
node end-point) to its tail-end (egress node end-point).
With end-to-end recovery, working LSPs are assumed to be resource
(link/node/SRLG) disjoint in the network so that they do not share
any failure probability, but this is not mandatory. With respect to
a given set of network resources, a pair of working/protecting LSPs
SHOULD be resource disjoint in case of dedicated recovery type (see
below). On the other hand, in case of shared recovery (see below), a
group of working LSPs SHOULD be mutually resource-disjoint in order
to allow for a (single and commonly) shared protecting LSP itself
resource-disjoint from each of the working LSPs. Note that resource
disjointness is a necessary (but not a sufficient) condition to
ensure LSP recoverability.

The present document addresses four types of end-to-end LSP
recovery: 1) 1+1 (unidirectional/bi-directional) protection, 2) 1:N
(N >= 1) LSP protection with extra-traffic, 3) pre-planned LSP re-
routing without extra-traffic (including shared mesh), and 4) full
LSP re-routing.

1) The simplest notion of end-to-end LSP protection is 1+1
   unidirectional protection. Using this type of protection, a
   protecting LSP is signaled over a dedicated resource-disjoint
   alternate path to protect an associated working LSP. Normal
   traffic is simultaneously sent on both LSPs and a selector is
   used at the egress node to receive traffic from one of the LSPs.
   If a failure occurs along one of the LSPs, the egress node
   selects the traffic from the valid LSP. No coordination is
   required between the end nodes when a failure/switchover occurs.

   In 1+1 bi-directional protection, a protecting LSP is signaled
   over a dedicated resource-disjoint alternate path to protect the
   working LSP. Normal traffic is simultaneously sent on both LSPs
   (in both directions) and a selector is used at both
   ingress/egress nodes to receive traffic from the same LSP. This
   requires co-ordination between the end-nodes when switching to
   the protecting LSP.

2) In 1:N (N >= 1) protection with extra-traffic, the protecting LSP
   is a fully provisioned and resource-disjoint LSP from the N
   working LSPs, that allows for carrying extra-traffic. The N
   working LSPs MAY be mutually resource-disjoint. Coordination
   between end-nodes is required when switching from one of the
   working to the protecting LSP. As the protecting LSP is fully

provisioned, default operations during protection switching are
specified for a protecting LSP carrying extra-traffic, but this
is not mandatory. Note that M:N protection is out of scope of
this document (though mechanisms it defines may be extended to
cover it).

3) Pre-planned LSP re-routing (or restoration) relies on the
   establishment between the same pair of end-nodes of a working LSP
   and a protecting LSP that is link/node/SRLG disjoint from the
   working one. Here, the recovery resources for the protecting LSP
   are pre-reserved but explicit action is required to activate
   (i.e. commit resource allocation at the data plane) a specific
   protecting LSP instantiated during the (pre-)provisioning phase.
   Since the protecting LSP is not "active" (i.e. fully
   instantiated), it can not carry any extra-traffic. This does not
   mean that the corresponding resources can not used by other LSPs.
   Therefore, this mechanism protects against working LSP(s)
   failure(s) but requires activation of the protecting LSP after
   working LSP failure occurrence. This requires restoration
   signaling along the protecting path. "Shared-mesh" restoration
   can be seen as a particular case of pre-planned LSP re-routing
   that reduces the recovery resource requirements by allowing

multiple protecting LSPs to share common link and node resources.
The recovery resources are pre-reserved but explicit action is
required to activate (i.e. commit resource allocation at the data
plane) a specific protecting LSP instantiated during the (pre-)
provisioning phase. This procedure requires restoration
signaling along the protecting path.

Note that in both cases, bandwidth pre-reserved for a protecting
(but not activated) LSP, can be made available for carrying extra
traffic. LSPs for extra traffic (with lower holding priority than
the protecting LSP) can then be established using the bandwidth
pre-reserved for the protecting LSP. Also, any lower priority LSP
that use the pre-reserved resources for the protecting LSP(s)
must be preempted during the activation of the protecting LSP.

4) Full LSP re-routing (or restoration) switches normal traffic to
   an alternate LSP that is not even partially established until
   after the working LSP failure occurs. The new alternate route is
   selected at the LSP head-end node, it may reuse resources of the
   failed LSP at intermediate nodes and may include additional
   intermediate nodes and/or links.

Crankback signaling (see [CRANK]) and LSP segment recovery (see
[SEGREC]) are further detailed in dedicated companion documents.

3. Relationship to Fast Reroute (FRR)

There is no impact to RSVP-TE Fast Reroute (FRR) [RFC4090]
introduced by end-to-end GMPLS recovery i.e. it is possible to use
either method defined in FRR with end-to-end GMPLS recovery.

The objects used and/or newly introduced by end-to-end recovery will
be ignored by [RFC4090] conformant implementations, and FRR can
operate on a per LSP basis as defined in [RFC4090].

4. Definitions

4.1 LSP Identification

This section reviews terms previously defined in [RFC2205],
[RFC3209], and [RFC3473]. LSP tunnels are identified by a
combination of the SESSION and SENDER_TEMPLATE objects (see also
[RFC3209]). The relevant fields are as follows:

IPv4 (or IPv6) tunnel end point address

    IPv4 (or IPv6) address of the egress node for the tunnel.

Tunnel ID

    A 16-bit identifier used in the SESSION that remains constant
    over the life of the tunnel.

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Extended Tunnel ID

    A 32-bit (or 16-byte) identifier used in the SESSION that
    remains constant over the life of the tunnel. Normally set to
    all zeros. Ingress nodes that wish to narrow the scope of a
    SESSION to the ingress-egress pair MAY place their IPv4 (or
    IPv6) address here as a globally unique identifier.

IPv4 (or IPv6) tunnel sender address

    IPv4 (or IPv6) address for a sender node.

LSP ID

    A 16-bit identifier used in the SENDER_TEMPLATE and FILTER_SPEC

that can be changed to allow a sender to share resources with
itself.

The first three fields are carried in the SESSION object (Path and
Resv message) and constitute the basic identification of the LSP
tunnel.

The last two fields are carried in the SENDER_TEMPLATE (Path
message) and FILTER_SPEC objects (Resv message). The LSP ID is used
to differentiate LSPs that belong to the same LSP Tunnel (as
identified by its Tunnel ID).

4.2 Recovery Attributes

The recovery attributes include all the parameters that determine
the status of a LSP within the recovery scheme to which it is
associated. These attributes are part of the PROTECTION object
introduced in Section 14.

4.2.1 LSP Status

The following bits are used in determining resource allocation and
status of the LSP within the group of LSPs forming the protected
entity:

- S (Secondary) bit: enables distinction between primary and
  secondary LSPs. A primary LSP is a fully established LSP for
  which the resource allocation has been committed at the data plane
  (i.e. full cross-connection has been performed). Both working and
  protecting LSPs can be primary LSPs. A secondary LSP is an LSP
  that has been provisioned in the control plane only and for which
  resource selection MAY have been done but for which the resource
  allocation has not been committed at the data plane (for instance,
  no cross-connection has been performed). Therefore, a secondary
  LSP is not immediately available to carry any traffic (requiring
  thus additional signaling to be available). A secondary LSP can

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

only be a protecting LSP. The (data plane) resources allocated for
a secondary LSP MAY be used by other LSPs until the primary LSP
fails over to the secondary LSP.

- P (Protecting) bit: enables distinction between working and
  protecting LSPs. A working LSP must be a primary LSP whilst a
  protecting LSP can be either a primary or a secondary LSP. When
  protecting LSP(s) are associated with working LSP(s), one also
  refers to the latter as protected LSPs.

Note: The combination "secondary working" is not valid (only
protecting LSPs can be secondary LSPs). Working LSPs are always
primary LSPs (i.e. fully established) whilst primary LSPs can be
either working or protecting LSPs.

- O (Operational) bit: this bit is set when a protecting LSP is
  carrying the normal traffic after protection switching (i.e.
  applies only in case of dedicated LSP protection or LSP protection
  with extra-traffic, see Section 4.2.2).

In this document, the PROTECTION object uses as a basis the
PROTECTION object defined in [RFC3471] and [RFC3473] and defines
additional fields within it. The fields defined in [RFC3471] and
[RFC3473] are unchanged by this document.

4.2.2 LSP Recovery

The following classification is used to distinguish the LSP
Protection Type with which LSPs can be associated at end-nodes (a
distinct value is associated with each Protection Type in the
PROTECTION object, see Section 14):

- Full LSP Re-routing: set if a primary working LSP is dynamically
  recoverable using (non pre-planned) head-end re-routing.

- Pre-planned LSP Re-routing without Extra-traffic: set if a
  protecting LSP is a secondary LSP that allows sharing of the
  pre-reserved recovery resources between one or more than one
  <sender;receiver> pair. When the secondary LSPs resources are not
  pre-reserved for a single <sender;receiver> pair, this type is
  referred to as "shared mesh" recovery.

- LSP Protection with Extra-traffic: set if a protecting LSP is a
  dedicated primary LSP that allows for extra-traffic transport
  and thus precludes any sharing of the recovery resources between
  more than one <sender;receiver> pair. This type includes 1:N LSP
  protection with extra-traffic.
- Dedicated LSP Protection: set if a protecting LSP does not allow
  sharing of the recovery resources nor the transport of extra-
  traffic (implying in the present context, duplication of the
  signal over both working and protecting LSPs as in 1+1 dedicated
  protection). Note also that this document makes a distinction

J.P.Lang et al.   Expires February 2007   7

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

between 1+1 unidirectional and bi-directional dedicated LSP
protection.

For LSP protection, in particular when the data plane provides
automated protection switching capability (see for instance ITU-T
[G.841] Recommendation), a Notification (N) bit is defined in the
PROTECTION object. It allows for distinction between protection
switching signaling via the control plane or via the data plane.

Note: this document assumes that Protection Type values have end-to-
end significance and that the same value is sent over the protected
and the protecting path. In this context, shared-mesh for instance,
appears from the end-nodes perspective as being simply an LSP re-
routing without extra-traffic services. The net result of this is
that a single bit (the S bit alone) does not allow determining
whether resource allocation should be performed and this *with
respect to* the status of the LSP within the protected entity. The
introduction of the P bit solves this problem unambiguously. These
bits MUST be processed on a hop-by-hop basis (independently of the
LSP Protection Type context). This allows for an easier
implementation of reversion signaling (see Section 12) but also
facilitates the transparent delivery of protected services since any
intermediate node is not required to know the semantic associated
with the incoming LSP Protection Type value.

4.3 LSP Association

The ASSOCIATION object, introduced in Section 16, is used to
associate the working and protecting LSPs.

When used for signaling the working LSP, the Association ID of the
ASSOCIATION object (see Section 16) identifies the protecting LSP.
When used for signaling the protecting LSP, this field identifies
the LSP protected by the protecting LSP.

5. 1+1 Unidirectional Protection

One of the simplest notions of end-to-end LSP protection is 1+1
unidirectional protection.

Consider the following network topology:

```
        A---B---C---D
         \         /
          E---F---G
```
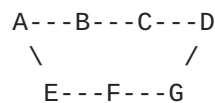
The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint,
ignoring the ingress/egress nodes A and D. A 1+1 protected path is
established from A to D over [A,B,C,D] and [A,E,F,G,D] and traffic
is transmitted simultaneously over both component paths (i.e. LSPs).

J.P.Lang et al.   Expires February 2007   8

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

During the provisioning phase, both LSPs are fully instantiated (and
thus activated) so that no resource sharing can be done along the
protecting LSP (nor can any extra-traffic be transported). It is
also RECOMMENDED to set the N bit since no protection switching
signaling is assumed in this case.

When a failure occurs (say at node B) and is detected at end-node D,
the receiver at D selects the normal traffic from the other LSP.
From this perspective, 1+1 unidirectional protection can be seen as
an uncoordinated protection switching mechanism acting independently
at both end-points. Also, for the LSP under failure condition, it is
RECOMMENDED to not set the Path_State_Removed Flag of the ERROR_SPEC
object (see [RFC3473]) upon PathErr message generation.

Note: it is necessary that both paths are SRLG disjoint to ensure
recoverability otherwise a single failure may impact both working
and protecting LSPs.

5.1. Identifiers

To simplify association operations, both LSPs belong to the same
session. Thus, the SESSION object MUST be the same for both LSPs.
The LSP ID, however, MUST be different to distinguish between the
two LSPs.

A new PROTECTION object (see Section 14) is included in the Path
message. This object carries the desired end-to-end LSP Protection
Type, in this case, "1+1 Unidirectional". This LSP Protection Type
value is applicable to both uni- and bi-directional LSPs.

To allow distinguishing the working LSP (from which the signal is
taken) from the protecting LSP, the working LSP is signaled by
setting in the PROTECTION object the S bit to 0, the P bit to 0, and
in the ASSOCIATION object, the Association ID to the protecting
LSP_ID. The protecting LSP is signaled by setting in the PROTECTION
object the S bit to 0, the P bit to 1, and in the ASSOCIATION
object, the Association ID to the associated protected LSP_ID.

After protection switching completes, and after reception of the
PathErr message, to keep track of the LSP from which the signal is
taken, the protecting LSP SHOULD be signaled with the O-bit set. The
formerly working LSP MAY be signaled with the A bit set in the
ADMIN_STATUS object (see [RFC3473]). This process assumes the tail-
end node has notified the head-end node that traffic selection
switchover has occurred.

[6](#). **1+1 Bi-directional Protection**

   1+1 bi-directional protection is a scheme that provides end-to-end
   protection for bi-directional LSPs.

J.P.Lang et al.   Expires February 2007   9

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

   Consider the following network topology:
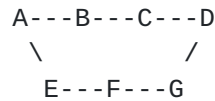
                         A---B---C---D
                          \   /
                           E---F---G

   The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint,
   ignoring the ingress/egress nodes A and D. A bi-directional LSP is
   established from A to D over each path and traffic is transmitted
   simultaneously over both LSPs. In this scheme, both end-points must
   receive traffic over the same LSP. Note also that both LSPs are
   fully instantiated (and thus activated) so that no resource sharing
   can be done along the protection path (nor can any extra-traffic be
   transported).

   When a failure is detected by one or both end-points of the LSP,
   both end-points must select traffic from the other LSP. This action
   must be coordinated between node A and D. From this perspective, 1+1
   bi-directional protection can be seen as a coordinated protection
   switching mechanism between both end-points.

   Note: it is necessary that both paths are SRLG disjoint to ensure
   recoverability, otherwise a single failure may impact both working
   and protecting LSPs.

[6.1](#). **Identifiers**

   To simplify association operations, both LSPs belong to the same
   session. Thus, the SESSION object MUST be the same for both LSPs.
   The LSP ID, however, MUST be different to distinguish between the
   two LSPs.

   A new PROTECTION object (see [Section 14](#)) is included in the Path
   message. This object carries the desired end-to-end LSP Protection
   Type, in this case, "1+1 Bi-directional". This LSP Protection Type
   value is only applicable to bi-directional LSPs.

   It is also desirable to allow distinguishing the working (LSP from
   which the signal is taken) from the protecting LSP. This is achieved
   for the working LSP by setting in the PROTECTION object the S bit to

0, the P bit to 0, and in the ASSOCIATION object, the Association ID
to the protecting LSP_ID. The protecting LSP is signaled by setting
in the PROTECTION object the S bit to 0, the P bit to 1 and in the
ASSOCIATION object the Association ID to the associated protected
LSP_ID.

6.2. End-to-End Switchover Request/Response

To co-ordinate the switchover between end-points, an end-to-end
switchover request/response exchange is needed since a failure
affecting one the LSPs results in both end-points switching to the

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

other LSP (resulting in receiving traffic from the other LSP) in
their respective directions.

The procedure is as follows:

   1. If an end-node (A or D) detects the failure of the working
      LSP (or a degradation of signal quality over the working
      LSP) or receives a Notify message including its SESSION
      object within the <upstream/downstream session list> (see
      [RFC3473]), and the new error code/sub-code "Notify Error/
      LSP Locally Failed" in the (IF_ID)_ERROR_SPEC object, it
      MUST begin receiving on the protecting LSP. Note that the
      <sender descriptor> or <flow descriptor> is also present in
      the Notify message that resolves any ambiguity and race
      condition since identifying (together with the SESSION
      object) the LSP under failure condition.

      This node MUST reliably send a Notify message including the
      MESSAGE_ID object to the other end-node (D or A,
      respectively) with the new error code/sub-code "Notify
      Error/LSP Failure" (Switchover Request) indicating the
      failure of the working LSP. This Notify message MUST be sent
      with the ACK_Desired flag set in the MESSAGE_ID object to
      request the receiver to send an acknowledgment for the
      message (see [RFC2961]).

      This (switchover request) Notify message MAY indicate the
      identity of the failed link or any other relevant
      information using the IF_ID ERROR_SPEC object (see
      [RFC3473]). In this case, the IF_ID ERROR_SPEC object
      replaces the ERROR_SPEC object in the Notify message,
      otherwise the corresponding (data plane) information SHOULD
      be received in the PathErr/ResvErr message.

2. Upon receipt of the (switchover request) Notify message, the
   end-node (D or A, respectively) MUST begin receiving from
   the protecting LSP.

   This node MUST reliably send a Notify message including the
   MESSAGE_ID object to the other end-node (A or D,
   respectively). This (switchover response) Notify message
   MUST also include a MESSAGE_ID_ACK object to acknowledge
   reception of the (switchover request) Notify message.

   This (switchover response) Notify message MAY indicate the
   identity of the failed link or any other relevant
   information using the IF_ID ERROR_SPEC object (see
   [RFC3473]).

   Note: upon receipt of the (switchover response) Notify
   message, the end-node (A or D, respectively) MUST send an
   Ack message to the other end-node to acknowledge its

   reception.

Since the intermediate nodes (B,C,E,F and G) are assumed to be GMPLS
RSVP-TE signaling capable, each node adjacent to the failure MAY
generate a Notify message directed either to the LSP head-end
(upstream direction) or the LSP tail-end (downstream direction) or
even both. Therefore, it is expected that these LSP terminating
nodes (that MAY also detect the failure of the LSP from the data
plane) provide either the right correlation mechanism to avoid
repetition of the above procedure or just discard subsequent Notify
messages corresponding to the same Session. In addition, for the LSP
under failure condition, it is RECOMMENDED to not set the
Path_State_ Removed Flag of the ERROR_SPEC object (see [RFC3473])
upon PathErr message generation.

After protection switching completes (step 2), and after reception
of the PathErr message, to keep track of the LSP from which the
signal is taken, the protecting LSP SHOULD be signaled with the O-
bit set. The formerly working LSP MAY be signaled with the A bit set
in the ADMIN_STATUS object (see [RFC3473]).

Note: when the N bit is set, the end-to-end switchover request/
response exchange described above only provides control plane
coordination (no actions are triggered at the data plane level).

7. 1:1 Protection with Extra-Traffic

The most common case of end-to-end 1:N protection is to establish,
between the same end-points, an end-to-end working LSP (thus, N = 1)
and a dedicated end-to-end protecting LSP that are mutually link/
node/SRLG disjoint. This protects against working LSP failure(s).

The protecting LSP is used for switchover when the working LSP
fails. GMPLS RSVP-TE signaling allows for the pre-provisioning of
protecting LSPs by indicating in the Path message (in the PROTECTION
object, see Section 14) that the LSPs are of type protecting. Here,
working and protecting LSPs are signaled as primary LSPs; both are
fully instantiated during the provisioning phase.

Although the resources for the protecting LSP are pre-allocated,
preemptable traffic may be carried end-to-end using this LSP. Thus,
the protecting LSP is capable of carrying extra-traffic with the
caveat that this traffic will be preempted if the working LSP fails.

The setup of the working LSP SHOULD indicate that the LSP head-end
and tail-end node wish to receive Notify messages using the NOTIFY
REQUEST object. The node upstream to the failure (upstream in terms
of the direction an Path message traverses) SHOULD send a Notify
message to the LSP head-end node, and the node downstream to the
failure SHOULD send an Notify message to the LSP tail-end node. Upon
receipt of the Notify messages, both the end-nodes MUST switch the
(normal) traffic from the working LSP to the pre-configured

protecting LSP (see Section 7.2). Moreover some coordination is
required if extra-traffic is carried over the end-to-end protecting
LSP. Note that if the working and the protecting LSP are established
between the same end-nodes no further notification is required to
indicate that the working LSPs are no longer protected.

Consider the following topology:

```
              A---B---C---D
               \         /
                E---F---G
```

The working LSP [A,B,C,D] could be protected by the protecting LSP
[A,E,F,G,D]. Both LSPs are fully instantiated (resources are
allocated for both working and protecting LSPs) and no resource
sharing can be done along the protection path since the primary
protecting LSP can carry extra-traffic.

Note: it is necessary that both paths are SRLG disjoint to ensure

recoverability otherwise a single failure may impact both working
and protecting LSPs.

7.1 Identifiers

To simplify association operations, both LSPs belong to the same
session. Thus, the SESSION object MUST be the same for both LSPs.
The LSP ID, however, MUST be different to distinguish between the
protected LSP carrying working traffic and the protecting LSP that
can carry extra-traffic.

A new PROTECTION object (see Section 14) is included in the Path
message used to setup the two LSPs. This object carries the desired
end-to-end LSP Protection Type, in this case, "1:N Protection with
Extra-Traffic". This LSP Protection Type value is applicable to both
uni- and bi-directional LSPs.

The working LSP is signaled by setting in the new PROTECTION object
the S bit to 0, the P bit to 0 and in the ASSOCIATION object the
Association ID to the protecting LSP_ID.

The protecting LSP is signaled by setting in the new PROTECTION
object the S bit to 0, the P bit to 1, and in the ASSOCIATION object
the Association ID to the associated protected LSP_ID.

7.2 End-to-End Switchover Request/Response

To co-ordinate the switchover between end-points, an end-to-end
switchover request/response is needed such that the affected LSP is
moved to the protecting LSP. Protection switching from the working
to the protecting LSP (implying preemption of extra-traffic carried
over the protecting LSP) must be initiated by one of the end-nodes
(A or D).

The procedure is as follows:

   1. If an end-node (A or D) detects the failure of the working
      LSP (or a degradation of signal quality over the working
      LSP) or receives a Notify message including its SESSION
      object within the <upstream/downstream session list> (see
      [RFC3473]), and the new error code/sub-code "Notify
      Error/LSP Locally Failed" in the (IF_ID)_ERROR_SPEC object,
      it disconnects the extra-traffic from the protecting LSP.
      Note that the <sender descriptor> or <flow descriptor> is
      also present in the Notify message that resolves any

ambiguity and race condition since identifying (together
with the SESSION object) the LSP under failure condition.

This node MUST reliably send a Notify message including the
MESSAGE_ID object to the other end-node (D or A,
respectively) with the new error code/sub-code "Notify
Error/LSP Failure" (Switchover Request) indicating the
failure of the working LSP. This Notify message MUST be sent
with the ACK_Desired flag set in the MESSAGE_ID object to
request the receiver to send an acknowledgment for the
message (see [RFC2961]).

This (switchover request) Notify message MAY indicate the
identity of the failed link or any other relevant
information using the IF_ID ERROR_SPEC object (see
[RFC3473]). In this case, the IF_ID ERROR_SPEC object
replaces the ERROR_SPEC object in the Notify message,
otherwise the corresponding (data plane) information SHOULD
be received in the PathErr/ResvErr message.

2. Upon receipt of the (switchover request) Notify message, the
   end-node (D or A, respectively) MUST disconnect the extra-
   traffic from the protecting LSP and begin sending/receiving
   normal traffic out/from the protecting LSP.

   This node MUST reliably send a Notify message including the
   MESSAGE_ID object to the other end-node (A or D,
   respectively). This (switchover response) Notify message
   MUST also include a MESSAGE_ID_ACK object to acknowledge
   reception of the (switchover request) Notify message.

   This (switchover response) Notify message MAY indicate the
   identity of the failed link or any other relevant
   information using the IF_ID ERROR_SPEC object (see
   [RFC3473]).

   Note: since the Notify message generated by the other end-
   node (A or D, respectively) is distinguishable from the one
   generated by an intermediate node, there is no possibility
   of connecting the extra traffic to the working LSP due to

the receipt of Notify message from an intermediate node.

3. Upon receipt of the (switchover response) Notify message,
   the end-node (A or D, respectively) MUST begin
   receiving/sending normal traffic from/out the protecting

LSP.

This node MUST also send an Ack message to the other end-
node (D or A, respectively) to acknowledge the reception of
the (switchover response) Notify message.

Note 1: a 2-phase protection switching signaling is used in the
present context, a 3-phase signaling (see [RFC4426]) that would
imply a notification message, a switchover request, and a switchover
response messages is not considered here. Also, when the protecting
LSPs do not carry extra-traffic, protection switching signaling as
defined in Section 6.2 MAY be used instead of the procedure
described in this section.

Note 2: when the N bit is set, the above end-to-end switchover
request/response exchange does only provide control plane
coordination (no actions are triggered at the data plane level).

After protection switching completes (step 3), and after reception
of the PathErr message, to keep track of the LSP from which the
normal traffic is taken, the protecting LSP SHOULD be signaled with
the O-bit set. In addition, the formerly working LSP MAY be signaled
with the A bit set in the ADMIN_STATUS object (see [RFC3473]).

7.3 1:N (N > 1) Protection with Extra-Traffic

1:N (N > 1) protection with extra-traffic assumes that the fully
provisioned protecting LSP is resource-disjoint from the N working
LSPs. This protecting LSP allows thus for carrying extra-traffic.
Note that the N working LSPs and the protecting LSP are all between
the same pair of end-points. In addition, the N working LSPs
(considered as identical in terms of traffic parameters) MAY be
mutually resource-disjoint. Coordination between end-nodes is
required when switching from one of the working to the protecting
LSP.

Each working LSP is signaled with both S bit and P bit set to 0. The
LSP Protection Type is set to 0x04 (1:N Protection with Extra-
Traffic) during LSP setup. Each Association ID points to the
protecting LSP ID.

The protecting LSP (carrying extra-traffic) is signaled with the S
bit set to 0 and the P bit set to 1. The LSP Protection Type is set
to 0x04 (1:N Protection with Extra-Traffic) during LSP setup. The
Association ID MUST be set by default to the LSP ID of the protected
LSP corresponding to N = 1.

J.P.Lang et al.   Expires February 2007   15

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Any signaling procedure applicable to 1:1 protection with extra-traffic equally applies to 1:N protection with extra-traffic.

8. Re-routing without Extra-Traffic

End-to-end (pre-planned) re-routing without extra-traffic relies on the establishment between the same pair of end-nodes of a working LSP and a protecting LSP that is link/node/SRLG disjoint from the working LSP. However, in this case the protecting LSP is not fully instantiated, thus, it can not carry any extra-traffic (note that this does not mean that the corresponding resources can not be used by other LSPs). Therefore, this mechanism protects against working LSP failure(s) but requires activation of the protecting LSP after failure occurrence.

Signaling is performed by indicating in the Path message (in the PROTECTION object, see Section 14) that the LSPs are of type working and protecting, respectively. Protecting LSPs are used for fast switchover when working LSPs fail. In this case, working and protecting LSPs are signaled as primary LSP and secondary LSP, respectively. Thus, only the working LSP is fully instantiated during the provisioning phase and for the protecting LSPs, no resources are committed at the data plane level (they are pre-reserved at the control plane level only). The setup of the working LSP SHOULD indicate (using the NOTIFY REQUEST object as specified in Section 4 of [RFC3473]) that the LSP head-end node (and possibly the tail-end node) wish to receive a Notify message upon LSP failure occurrence. Upon receipt of the Notify message, the head-end node MUST switch the (normal) traffic from the working LSP to the protecting LSP after its activation. Note that since the working and the protecting LSP are established between the same end-nodes no further notification is required to indicate that the working LSPs are no longer protected.

To make bandwidth pre-reserved for a protecting (but not activated) LSP, available for extra traffic this bandwidth could be included in the advertised Unreserved Bandwidth at priority lower (means numerically higher) than the Holding Priority of the protecting LSP. In addition, the Max LSP Bandwidth field in the Interface Switching Capability Descriptor sub-TLV should reflect the fact that the bandwidth pre-reserved for the protecting LSP is available for extra traffic. LSPs for extra traffic then can be established using the bandwidth pre-reserved for the protecting LSP by setting (in the Path message) the Setup Priority field of the SESSION_ATTRIBUTE object to X (where X is the Setup Priority of the protecting LSP) and the Holding Priority field at least to X+1. Also, if the resources pre-reserved for the protecting LSP are used by lower priority LSPs, these LSPs MUST be preempted when the protecting LSP is activated (see Section 10).

Consider the following topology:

```
              A---B---C---D
               \         /
                E---F---G
```
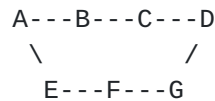
The working LSP [A,B,C,D] could be protected by the protecting LSP
[A,E,F,G,D]. Only the protected LSP is fully instantiated (resources
are only allocated for the working LSP). Therefore, the protecting
LSP can not carry any extra-traffic. When a failure is detected on
the working LSP (say at B), the error is propagated and/or notified
(using a Notify message with the new error code/sub-code "Notify
Error/LSP Locally Failed" in the (IF_ID)_ERROR_SPEC object) to the
ingress node (A). Upon reception, the latter activates the secondary
protecting LSP instantiated during the (pre-)provisioning phase.
This requires:
(1) the ability to identify a "secondary protecting LSP" (hereby
     called the "secondary LSP") used to recover another primary
     working LSP (hereby called the "protected LSP")
(2) the ability to associate the secondary LSP with the protected
     LSP
(3) the capability to activate a secondary LSP after failure
     occurrence.

In the following subsections, these features are described in more
detail.

8.1 Identifiers

To simplify association operations, both LSPs (i.e. the protected
and the secondary LSPs) belong to the same session. Thus, the
SESSION object MUST be the same for both LSPs. The LSP ID, however,
MUST be different to distinguish between the protected LSP carrying
working traffic and the secondary LSP that can not carry extra-
traffic.

A new PROTECTION object (see Section 14) is used to setup the two
LSPs. This object carries the desired end-to-end LSP Protection Type
(in this case, "Re-routing without Extra-Traffic"). This LSP
Protection Type value is applicable to both uni- and bi-directional
LSPs.

8.2 Signaling Primary LSPs

The new PROTECTION object is included in the Path message during signaling of the primary working LSP, with the end-to-end LSP Protection Type value set to "Re-routing without Extra-Traffic".

Primary working LSPs are signaled by setting in the new PROTECTION object the S bit to 0, the P bit to 0 and in the ASSOCIATION object the Association ID to the associated secondary protecting LSP_ID.

8.3 Signaling Secondary LSPs

The new PROTECTION object is included in the Path message during signaling of secondary protecting LSPs, with the end-to-end LSP Protection Type value set to "Re-routing without Extra-Traffic".

Secondary protecting LSPs are signaled by setting in the new PROTECTION object the S bit and the P bit to 1 and in the ASSOCIATION object the Association ID to the associated primary working LSP_ID, which MUST be known before signaling of the secondary LSP.

With this setting, the resources for the secondary LSP SHOULD be pre-reserved, but not committed at the data plane level meaning that the internals of the switch need not be established until explicit action is taken to activate this secondary LSP. Activation of a secondary LSP is done using a modified Path message with the S bit set to 0 in the PROTECTION object. At this point, the link and node resources must be allocated for this LSP that becomes a primary LSP (ready to carry normal traffic).

From [RFC3945], the secondary LSP is setup with resource pre-reservation but with or without label pre-selection (both allowing sharing of the recovery resources). In the former case (defined as the default), label allocation during secondary LSP signaling does not require any specific procedure compared to [RFC3473]. However, in the latter case, label (and thus resource) re-allocation MAY occur during the secondary LSP activation. This means that during the LSP activation phase, labels MAY be re-assigned (with higher precedence over existing label assignment, see also [RFC3471]).

Note: under certain circumstances (e.g. when pre-reserved protecting resources are used by lower priority LSPs), it MAY be desirable to perform the activation of the secondary LSP in the upstream direction (Resv trigger message) instead of using the default downstream activation. In this case, any mis-ordering and any mis-interpretation between a refresh Resv (along the lower priority LSP)

and a trigger Resv message (along the secondary LSP) MUST be avoided
at any intermediate node. For this purpose, upon reception of the
Path message, the egress node MAY include the PROTECTION object in
the Resv message. The latter is then processed on a hop by hop basis
to activate the secondary LSP until reaching the ingress node. The
PROTECTION object included in the Path message MUST be set as
specified in this Section. In this case, the PROTECTION object with
the S bit MUST be set to 0 and included in the Resv message sent in
the upstream direction. The upstream activation behavior SHOULD be
configurable on a local basis. Details concerning lower priority LSP
preemption upon secondary LSP activation are provided in Section 10.

9. Shared-Mesh Restoration

An approach to reduce recovery resource requirements is to have
protection LSPs sharing network resources when the working LSPs that
they protect are physically (i.e., link, node, SRLG, etc.) disjoint.

This mechanism is referred to as shared mesh restoration and is
described in [RFC4426]. Shared-mesh restoration can be seen as a
particular case of pre-planned LSP re-routing (see Section 8) that
reduces the recovery resource requirements by allowing multiple
protecting LSPs to share common link and node resources. Here also,
the recovery resources for the protecting LSPs are pre-reserved
during the provisioning phase, thus an explicit signaling action is
required to activate (i.e. commit resource allocation at the data
plane) a specific protecting LSP instantiated during the (pre-)
provisioning phase. This requires restoration signaling along the
protecting LSP.

To make bandwidth pre-reserved for a protecting (but not activated)
LSP, available for extra traffic this bandwidth could be included in
the advertised Unreserved Bandwidth at priority lower (means
numerically higher) than the Holding Priority of the protecting LSP.
In addition, the Max LSP Bandwidth field in the Interface Switching
Capability Descriptor sub-TLV should reflect the fact that the
bandwidth pre-reserved for the protecting LSP is available for extra
traffic. LSPs for extra traffic then can be established using the
bandwidth pre-reserved for the protecting LSP by setting (in the
Path message) the Setup Priority field of the SESSION_ATTRIBUTE
object to X (where X is the Setup Priority of the protecting LSP)
and the Holding Priority field at least to X+1. Also, if the
resources pre-reserved for the protecting LSP are used by lower
priority LSPs, these LSPs MUST be preempted when the protecting LSP
is activated (see Section 10). Further, if the recovery resources

are shared between multiple protecting LSPs, the corresponding
working LSPs head-end nodes must be informed that they are no longer
protected when the protecting LSP is activated to recover the normal
traffic for the working LSP under failure.

Consider the following topology:


```
                     A---B---C---D
                      \         /
                       E---F---G
                      /         \
                     H---I---J---K
```
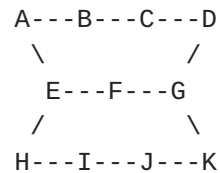

The working LSPs [A,B,C,D] and [H,I,J,K] could be protected by
[A,E,F,G,D] and [H,E,F,G,K], respectively. Per [RFC3209], in order
to achieve resource sharing during the signaling of these protecting
LSPs, they must have the same Tunnel Endpoint Address (as part of
their SESSION object). However, these addresses are not the same in
this example. Resource sharing along E, F, G can only be achieved if
the nodes E, F and G recognize that the LSP Protection Type of the
secondary LSPs is set to "Re-routing without Extra-Traffic" (see
PROTECTION object, Section 14) and acts accordingly. In this case,

the protecting LSPs are not merged (which is useful since the paths
diverge at G), but the resources along E, F, G can be shared.

When a failure is detected on one of the working LSPs (say at B),
the error is propagated and/or notified (using a Notify message with
the new error code/sub-code "Notify Error/LSP Locally Failed" in the
(IF_ID)_ERROR_SPEC object) to the ingress node (A). Upon reception,
the latter activates the secondary protecting LSP (see Section 8).
At this point, it is important that a failure on the other LSP (say
at J) does not cause the other ingress (H) to send the data down the
protecting LSP since the resources are already in use. This can be
achieved by node E using the following procedure. When the capacity
is first reserved for the protecting LSP, E should verify that the
LSPs being protected ([A,B,C,D] and [H,I,J,K], respectively) do not
share any common resources. Then, when a failure occurs (say at B)
and the protecting LSP [A,E,F,G,D] is activated, E should notify H
that the resources for the protecting LSP [H,E,F,G,K] are no longer
available.

The following sub-sections details how shared mesh restoration can

be implemented in an interoperable fashion using GMPLS RSVP-TE
extensions (see [RFC3473]). This includes:
(1) the ability to identify a "secondary protecting LSP" (hereby
     called the "secondary LSP") used to recover another primary
     working LSP (hereby called the "protected LSP")
(2) the ability to associate the secondary LSP with the protected
     LSP
(3) the capability to include information about the resources used
     by the protected LSP while instantiating the secondary LSP.
(4) the capability to instantiate during the provisioning phase
     several secondary LSPs in an efficient manner.
(5) the capability to activate a secondary LSP after failure
     occurrence.

   In the following subsections, these features are described in
   detail.

9.1. Identifiers

   To simplify association operations, both LSPs (i.e. the protected
   and the secondary LSPs) belong to the same session. Thus, the
   SESSION object MUST be the same for both LSPs. The LSP ID, however,
   MUST be different to distinguish between the protected LSP carrying
   working traffic and the secondary LSP that can not carry extra-
   traffic.

   A new PROTECTION object (see Section 14) is used to setup the two
   LSPs. This object carries the desired end-to-end LSP Protection
   Type, in this case, "Re-routing without Extra-Traffic". This LSP
   Protection Type value is applicable to both uni- and bi-directional
   LSPs.

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

9.2 Signaling Primary LSPs

   The new PROTECTION object is included in the Path message during
   signaling of the primary working LSPs, with the end-to-end LSP
   Protection Type value set to "Re-routing without Extra-Traffic".

   Primary working LSPs are signaled by setting in the new PROTECTION
   object the S bit to 0, the P bit to 0 and in the ASSOCIATION object
   the Association ID to the associated secondary protecting LSP_ID.

9.3 Signaling Secondary LSPs

   The new PROTECTION object is included in the Path message during

signaling of the secondary protecting LSPs, with the end-to-end LSP
Protection Type value set to "Re-routing without Extra-Traffic".

Secondary protecting LSPs are signaled by setting in the new
PROTECTION object the S bit and the P bit to 1 and in the
ASSOCIATION object the Association ID to the associated primary
working LSP_ID, which MUST be known before signaling of the
secondary LSP. Moreover, the Path message used to instantiate the
secondary LSP SHOULD include at least one PRIMARY PATH ROUTE object
(see Section 15) that further allows for recovery resource sharing
at each intermediate node along the secondary path.

With this setting, the resources for the secondary LSP SHOULD be
pre-reserved, but not committed at the data plane level meaning that
the internals of the switch need not be established until explicit
action is taken to activate this LSP. Activation of a secondary LSP
is done using a modified Path message with the S bit set to 0 in the
PROTECTION object. At this point, the link and node resources must
be allocated for this LSP that becomes a primary LSP (ready to carry
normal traffic).

From [RFC3945], the secondary LSP is setup with resource pre-
reservation but with or without label pre-selection (both allowing
sharing of the recovery resources). In the former case (defined as
the default), label allocation during secondary LSP signaling does
not require any specific procedure compared to [RFC3473]. However,
in the latter case, label (and thus resource) re-allocation MAY
occur during the secondary LSP activation. This means that during
the LSP activation phase, labels MAY be re-assigned (with higher
precedence over existing label assignment, see also [RFC3471]).

10. LSP Preemption

When protecting resources are only pre-reserved for the secondary
LSPs, they MAY be used to setup lower priority LSPs. In this case,
these resources MUST be preempted when the protecting LSP is
activated. An additional condition raises from mis-connection
avoidance between the secondary protecting LSP being activated and
the low priority LSP(s) being preempted. Procedure to be applied

when the secondary protecting LSP (i.e. the pre-empting LSP) Path
message reaches a node using the resources for lower priority LSP(s)
(i.e. pre-empted LSP(s)) is as follows:

1. Deallocate resources to be used by the pre-empting LSP and
release the cross-connection. Note that if the pre-empting LSP is

bi-directional, these resources may come from one or two lower
priority LSPs, and if from two LSPs, they may be uni- or bi-
directional. The pre-empting node SHOULD NOT send the Path message
before the deallocation of resources has completed since this may
lead to the downstream path becoming misconnected if the downstream
node is able to re-assign the resources more quickly.

2. Send PathTear and PathErr messages with the new error code/sub-
code "Policy Control failure/Hard Pre-empted" and the Path_State_
Removed flag set for the pre-empted LSP(s).

3. Reserve the pre-empted resources for the protecting LSP. The pre-
empting node MUST NOT cross-connect the upstream resources of a bi-
directional pre-empting LSP.

4. Send the Path message.

5. Upon reception of a trigger Resv message from the downstream
node, cross-connect the downstream path resources and if the pre-
empting LSP is bi-directional, perform cross-connection for the
upstream path resources.

Note that step 1 may cause alarms to be raised for the pre-empted
LSP. If alarm suppression is desired the pre-empting node MAY insert
the following steps before step 1.

1a. Before deallocating resources send a Resv message including an
    ADMIN_STATUS object to disable alarms for the pre-empted LSP.
1b. Receive a Path message indicating that alarms are disabled.

At the downstream node (with respect to the pre-empting LSP) the
processing is RECOMMENDED to be as follows:

1. Receive PathTear (and/or PathErr) message for the pre-empted
   LSP(s).

2a.Release the resources associated with the LSP on the interface
   to the pre-empting LSP, remove any cross-connection and release
   all other resources associated with the pre-empted LSP.
2b.Forward the PathTear (and/or PathErr) message per [RFC3473].

3. Receive the Path message for the pre-empting LSP and process as
   normal, forwarding it to the downstream node.

4. Receive the Resv message for the pre-empting LSP and process as
   normal, forwarding it to the upstream node.

J.P.Lang et al.   Expires February 2007   22

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

LSP re-routing, on the other hand, switches normal traffic to an
alternate LSP that is fully established only after failure
occurrence. The new (alternate) route is selected at the LSP head-
end and may reuse intermediate nodes included in the original route;
it may also include additional intermediate nodes. For strict-hop
routing, TE requirements can be directly applied to the route
computation, and the failed node or link can be avoided. However, if
the failure occurred within a loose-routed hop, the head-end node
may not have enough information to reroute the LSP around the
failure. Crankback signaling (see [[CRANK](#)]) and route exclusion
techniques (see [[XRO](#)]) MAY be used in this case.

The alternate route MAY be either computed on demand (that is, when
the failure occurs; this is referred to as full LSP re-routing) or
pre-computed and stored for use when the failure is reported. The
latter offers faster restoration time. There is, however, a risk
that the alternate route will become out of date through other
changes in the network - this can be mitigated to some extent by
periodic recalculation of idle alternate routes.

(Full) LSP re-routing will be initiated by the head-end node that
has either detected the LSP failure or received a Notify message
and/or a PathErr message with the new error code/sub-code "Notify
Error/LSP Locally Failed" for this LSP. The new LSP resources can be
established using the make-before-break mechanism, where the new LSP
is setup before the old LSP is torn down. This is done by using the
mechanisms of the SESSION_ATTRIBUTE object and the Shared-Explicit
(SE) reservation style (see [[RFC3209](#)]). Both the new and old LSPs
can share resources at common nodes.

Note that the make-before-break mechanism is not used to avoid
disruption to the normal traffic flow (the latter has already been
broken by the failure that is being repaired). However, it is
valuable to retain the resources allocated on the original LSP that
will be re-used by the new alternate LSP.

[11.1](#) Identifiers

The Tunnel End Point Address, Tunnel ID, Extended Tunnel ID, Tunnel
Sender Address uniquely identify both the old and new LSPs. Only the
LSP_ID value differentiates the old from the new alternate LSP. The
new alternate LSP is setup before the old LSP is torn down using
Shared-Explicit (SE) reservation style. This ensures that the new
(alternate) LSP is established without double counting resource
requirements along common segments.

The alternate LSP MAY be setup before any failure occurrence with SE

style resource reservation, the latter shares the same Tunnel End
Point Address, Tunnel ID, Extended Tunnel ID, and Tunnel Sender

J.P.Lang et al.   Expires February 2007   23

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Address with the original LSP (i.e. only the LSP ID value MUST be
different).

In both cases, the Association ID of the ASSOCIATION object MUST be
set to the LSP ID value of the signaled LSP.

11.2 Signaling Re-routable LSPs

A new PROTECTION object is included in the Path message during
signaling of dynamically re-routable LSPs, with the end-to-end LSP
Protection Type value set to "Full Re-routing". These LSPs that can
be either uni- or bi-directional are signaled by setting in the
PROTECTION object the S bit to 0, the P bit to 0 and the Association
ID value to the LSP_ID value of the signaled LSP. Any specific
action to be taken during the provisioning phase is up to the end-
node local policy.

Note: when the end-to-end LSP Protection Type is set to
"Unprotected", both S and P bit MUST be set to 0 and the LSP SHOULD
NOT be re-routed at the head-end node after failure occurrence. The
Association_ID value MUST be set to the LSP_ID value of the signaled
LSP. This does not mean that the Unprotected LSP can not be re-
established for other reasons such as path re-optimization and
bandwidth adjustment driven by policy conditions.

12. Reversion

Reversion refers to a recovery switching operation, where the normal
traffic returns to (or remains on) the working LSP when it has
recovered from the failure. Reversion implies that resources remain
allocated to the LSP that was originally routed over them even after
a failure. It is important to have mechanisms that allow reversion
to be performed with minimal service disruption and reconfiguration.

For "1+1 bi-directional Protection", reversion to the recovered LSP
occurs by using the following sequence:

1. Clear the A bit of the ADMIN_STATUS object if set for the
   recovered LSP.

2. Then, apply the method described here below to switch normal
   traffic back from the protecting to the recovered LSP. This is
   performed by using the new error code/sub-code "Notify Error/LSP

Recovered" (Switchback Request).

   The procedure is as follows:

   1. The initiating (source) node sends the normal traffic onto
      both the working and the protecting LSPs. Once completed, the
      source node sends reliably a Notify message to the destination
      with the new error code/sub-code "Notify Error/LSP Recovered"
      (Switchback Request). This Notify message includes the

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

      MESSAGE_ID object. The ACK_Desired flag MUST be set in this
      object to request the receiver to send an acknowledgment for
      the message (see [RFC2961]).

   2. Upon receipt of this message, the destination selects the
      traffic from the working LSP. At the same time, it transmits
      the traffic onto both the working and protecting LSP.

      The destination then sends reliably a Notify message to the
      source confirming the completion of the operation. This
      message includes the MESSAGE_ID_ACK object to acknowledge
      reception of the received Notify message. This Notify message
      also includes the MESSAGE_ID object. The ACK_Desired flag MUST
      be set in this object to request the receiver to send an
      acknowledgment for the message (see [RFC2961]).

   3. When the source node receives this Notify message, it switches
      to receive traffic from the working LSP.

      The source node then sends an Ack message to the destination
      node confirming that the LSP has been reverted.

   3. Finally, clear the O bit of the PROTECTION object sent over the
      protecting LSP.

   For "1:N Protection with Extra-traffic", reversion to the recovered
   LSP occurs by using the following sequence:

   1. Clear the A bit of the ADMIN_STATUS object if set for the
      recovered LSP.

   2. Then, apply the method described here below to switch normal
      traffic back from the protecting to the recovered LSP. This is
      performed by using the new error code/sub-code "Notify Error/LSP
      Recovered" (Switchback Request).

The procedure is as follows:

1. The initiating (source) node sends the normal traffic onto both the working and the protecting LSPs. Once completed, the source node sends reliably a Notify message to the destination with the new error code/sub-code "Notify Error/LSP Recovered" (Switchback Request). This Notify message includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

2. Upon receipt of this message, the destination selects the traffic from the working LSP. At the same time, it transmits the traffic onto both the working and protecting LSP.

   The destination then sends reliably a Notify message to the

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

   source confirming the completion of the operation. This message includes the MESSAGE_ID_ACK object to acknowledge reception of the received Notify message. This Notify message also includes the MESSAGE_ID object. The ACK_Desired flag MUST be set in this object to request the receiver to send an acknowledgment for the message (see [RFC2961]).

3. When the source node receives this Notify message, it switches to receive traffic from the working LSP, and stops transmitting traffic on the protecting LSP.

   The source node then sends an Ack message to the destination node confirming that the LSP has been reverted.

4. Upon receipt of this message, the destination node stops transmitting traffic along the protecting LSP.

3. Finally, clear the O bit of the PROTECTION object sent over the protecting LSP.

For "Re-routing without Extra-traffic" (including the shared recovery case), reversion implies that the formerly working LSP has not been torn down by the head-end node upon PathErr message reception i.e. the head-end node kept refreshing the working LSP under failure condition. This ensures that the exact same resources are retrieved after reversion switching (except if the working LSP required re-signaling). Re-activation is performed using the following sequence:

1. Clear the A bit of the ADMIN_STATUS object if set for the
   recovered LSP.

2. Then, apply the method described here below to switch normal
   traffic back from the protecting to the recovered LSP. This is
   performed by using the new error code/sub-code "Notify Error/LSP
   Recovered" (Switchback Request).

   The procedure is as follows:

   1. The initiating (source) node sends the normal traffic onto
      both the working and the protecting LSPs. Once completed, the
      source node sends reliably a Notify message to the destination
      with the new error code/sub-code "Notify Error/LSP Recovered"
      (Switchback Request). This Notify message includes the
      MESSAGE_ID object. The ACK_Desired flag MUST be set in this
      object to request the receiver to send an acknowledgment for
      the message (see [RFC2961]).

   2. Upon receipt of this message, the destination selects the
      traffic from the working LSP. At the same time, it transmits
      the traffic onto both the working and protecting LSP.

      The destination then sends reliably a Notify message to the
      source confirming the completion of the operation. This
      message includes the MESSAGE_ID_ACK object to acknowledge
      reception of the received Notify message. This Notify message
      also includes the MESSAGE_ID object. The ACK_Desired flag MUST
      be set in this object to request the receiver to send an
      acknowledgment for the message (see [RFC2961]).

   3. When the source node receives this Notify message, it switches
      to receive traffic from the working LSP, and stops
      transmitting traffic on the protecting LSP.

      The source node then sends an Ack message to the destination
      node confirming that the LSP has been reverted.

   4. Upon receipt of this message, the destination node stops
      transmitting traffic along the protecting LSP.

3. Finally, de-activate the protecting LSP by setting the S bit to 1
   in the PROTECTION object sent over the protecting LSP.

13. Recovery Commands

This section specifies the control plane behavior when using several
commands (see [RFC4427]) that can be used to influence the recovery
operations.

A. Lockout of recovery LSP:

The Lockout bit (L bit) of the ADMIN_STATUS object is used following
the rules defined in Section 8 of [RFC3471] and Section 7 of
[RFC3473]. The L bit must be set together with the Reflect (R) bit
in the ADMIN_STATUS object sent in the Path message. Upon reception
of the Resv message with the L bit set, this forces the recovery LSP
to be temporarily unavailable to transport traffic (either normal or
extra traffic). Unlock is performed by clearing the L bit, following
the rules defined in Section 7 of [RFC3473]. This procedure is only
applicable when the LSP Protection Type Flag is set to either 0x04
(1:N Protection with Extra-Traffic), or 0x08 (1+1 Unidirectional
Protection) or 0x10 (1+1 Bi-directional Protection).

The updated format of the ADMIN_STATUS object to include the L bit
is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|          Length          | Class-Num(196)|   C-Type (1)  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|R|                   Reserved              |L|I|C|T|A|D|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

J.P.Lang et al.   Expires February 2007   27

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Lockout (L): 1 bit

    When set, indicates forces the recovery LSP to be temporarily
    unavailable to transport traffic (either normal or extra
    traffic).

The R (Reflect), T (Testing), A (Administratively down) and D
(Deletion in progress) bits are defined in [RFC3471]. The C (Call
control) bit is defined in [GMPLS-CALL], and the I (Inhibit alarm
communication) bit in [ALARM].

B. Lockout of normal traffic:

The O bit of the PROTECTION object is set to 1 to force the recovery
LSP to be temporarily unavailable to transport normal traffic. This

operation MUST NOT occur unless the working LSP is carrying the
normal traffic. Unlock is performed by clearing the O bit over the
protecting LSP. This procedure is only applicable when the LSP
Protection Type Flag is set to either 0x04 (1:N Protection with
Extra-Traffic), or 0x08 (1+1 Unidirectional Protection) or 0x10 (1+1
Bi-directional Protection).

C. Forced switch for normal traffic:

Recovery signaling is initiated that switches normal traffic to the
recovery LSP following the procedures defined in Section 6, 7, 8 and
9.

D. Requested switch for normal traffic:

Recovery signaling is initiated that switches normal traffic to the
recovery LSP following the procedures defined in Section 6, 7, 8 and
9. This, except if a fault condition exists on other LSPs/spans
(including the recovery LSP) or an equal or higher priority switch
command is in effect.

E. Requested switch for recovery LSP:

Recovery signaling is initiated that switches normal traffic to the
working LSP following the procedure defined in Section 12. This,
except if a fault condition exists on the working LSP or an equal or
higher priority switch command is in effect.

14. PROTECTION Object

This section describes the extensions to the PROTECTION object to
broaden its applicability to end-to-end LSP recovery.

14.1 Format

The format of the PROTECTION Object (Class-Num = 37, C-Type = 2,
suggested value, TBA by IANA) is as follows:

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |            Length             | Class-Num(37) | C-Type (TBA)  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |S|P|N|O| Reserved  | LSP Flags |     Reserved      | Link Flags|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

```
|                            Reserved                           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Secondary (S): 1 bit

   When set to 1, this bit indicates that the requested LSP is a
   secondary LSP. When set to 0 (default), it indicates that the
   requested LSP is a primary LSP.

Protecting (P): 1 bit

   When set to 1, this bit indicates that the requested LSP is a
   protecting LSP. When set to 0 (default), it indicates that the
   requested LSP is a working LSP. The combination, S set to 1
   with P set to 0 is not valid.

Notification (N): 1 bit

   When set to 1, this bit indicates that the control plane
   message exchange is only used for notification during
   protection switching. When set to 0 (default), it indicates
   that the control plane message exchanges are used for
   protection switching purposes. The N bit is only applicable
   when the LSP Protection Type Flag is set to either 0x04 (1:N
   Protection with Extra-Traffic), or 0x08 (1+1 Unidirectional
   Protection) or 0x10 (1+1 Bi-directional Protection). The N bit
   MUST be set to 0 in any other case.

Operational (O): 1 bit

   When set to 1, this bit indicates that the protecting LSP is
   carrying the normal traffic after protection switching. The O
   bit is only applicable when the P bit is set to 1 and the LSP
   Protection Type Flag is set to either 0x04 (1:N Protection
   with Extra-Traffic), or 0x08 (1+1 Unidirectional Protection)
   or 0x10 (1+1 Bi-directional Protection). The O bit MUST be set
   to 0 in any other case.

Reserved: 5 bits

   This field is reserved. It MUST be set to zero on transmission
   and MUST be ignored on receipt. These bits SHOULD be passed
   through unmodified by transit nodes.

LSP (Protection Type) Flags: 6 bits

Indicates the desired end-to-end LSP recovery type. A value of
0 implies that the LSP is "Unprotected". Only one value SHOULD
be set at a time. The following values are defined. All other
values are reserved.

```
0x00    Unprotected
0x01    (Full) Re-routing
0x02    Re-routing without Extra-Traffic
0x04    1:N Protection with Extra-Traffic
0x08    1+1 Unidirectional Protection
0x10    1+1 Bi-directional Protection
```

Reserved: 10 bits

This field is reserved. It MUST be set to zero on transmission
and MUST be ignored on receipt. These bits SHOULD be passed
through unmodified by transit nodes.

Link Flags: 6 bits

Indicates the desired link protection type (see [RFC3471]).

Reserved field: 32 bits

Encoding of this field is detailed in [SEGREC].

14.2 Processing

Intermediate and egress nodes processing a Path message containing a
PROTECTION object MUST verify that the requested LSP Protection Type
can be satisfied by the incoming interface. If it can not, the node
MUST generate a PathErr message, with the new error code/sub-code
"Routing problem/Unsupported LSP Protection".

Intermediate nodes processing a Path message containing a PROTECTION
object with the LSP Protection Type 0x02 (Re-routing without Extra-
Traffic) value set and a PRIMARY PATH ROUTE object (see Section 15)
MUST verify that the requested LSP Protection Type can be supported
by the outgoing interface. If it can not, the node MUST generate a
PathErr message with the new error code/sub-code "Routing
problem/Unsupported LSP Protection".

15. PRIMARY PATH ROUTE Object

The PRIMARY PATH ROUTE object (PPRO) is defined to inform nodes
along the path of a secondary protecting LSP about which resources
(link/nodes) are being used by the associated primary protected LSP
(as specified by the Association ID field). If the LSP Protection
Type value is set to 0x02 (Re-routing without Extra-Traffic), this
object SHOULD be present in the Path message for the pre-

J.P.Lang et al.   Expires February 2007   30

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

   provisioning of the secondary protecting LSP to enable recovery
   resource sharing between one or more secondary protecting LSPs (see
   Section 9). This document does not assume or preclude any other
   usage for this object.

   PRIMARY PATH ROUTE objects carry information extracted from the
   EXPLICIT ROUTE object and/or the RECORD ROUTE object of the primary
   working LSPs they protect. Selection of the PPRO content is up to
   local policy of the head-end node that initiates the request.
   Therefore, the information included in these objects can be used as
   policy-based admission control to ensure that recovery resources are
   only shared between secondary protecting LSPs whose associated
   primary LSPs have link/node/SRLG disjoint paths.

15.1 Format

   The primary path route is specified via the PRIMARY_PATH_ROUTE
   object (PPRO). The Primary Path Route Class Number (Class-Num) of
   form 0bbbbbbb is TBA by IANA.

   Currently one C-Type (Class-Type) is defined, Type 1, Primary Path
   Route. The PRIMARY_PATH_ROUTE object has the following format:

   Class-Num = TBA by IANA (of form 0bbbbbbb), C-Type = 1 (suggested)

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |                                                               |
    //                        (Subobjects)                        //
    |                                                               |
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

   The contents of a PRIMARY_PATH_ROUTE object are a series of
   variable-length data items called subobjects (see Section 15.3).

   To signal a secondary protecting LSP, the Path message MAY include
   one or multiple PRIMARY_PATH_ROUTE objects, where each object is
   meaningful. The latter is useful when a given secondary protecting
   LSP must be link/node/SRLG disjoint from more than one primary LSP
   (i.e. is protecting more than one primary LSP).

15.2 Subobjects

   The PRIMAY_PATH_ROUTE object is defined as a list of variable-length
   data items called subobjects. These subobjects are derived from the

subobjects of the EXPLICIT ROUTE and/or RECORD ROUTE object of the
primary working LSP(s).

Each subobject has its own length field. The length contains the
total length of the subobject in bytes, including the Type and

J.P.Lang et al.   Expires February 2007   31

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

Length fields. The length MUST always be a multiple of 4, and at
least 4.

The following subobjects are currently defined for the PRIMARY PATH
ROUTE object:

- Sub-Type 1: IPv4 Address (see [RFC3209])
- Sub-Type 2: IPv6 Address (see [RFC3209])
- Sub-Type 3: Label (see [RFC3473])
- Sub-Type 4: Unnumbered Interface (see [RFC3477])

An empty PPRO with no subobjects is considered as illegal. If there
is no first subobject, the corresponding Path message is also in
error and the receiving node SHOULD return a PathErr message with
the new error code/sub-code "Routing Problem/Bad PRIMARY PATH_ROUTE
object".

Note: an intermediate node processing a PPRO can derive SRLG
identifiers from the local IGP-TE database using its Type 1, 2 or 4
subobject values as pointers to the corresponding TE Links (assuming
each of them has an associated SRLG TE attribute).

15.3 Applicability

The PRIMARY_PATH_ROUTE object MAY only be used when all GMPLS nodes
along the path support the PRIMARY_PATH_ROUTE object and a secondary
protecting LSP is being requested. The PRIMARY_PATH_ROUTE object is
assigned a class value of the form 0bbbbbbb. Receiving GMPLS nodes
along the path that do not support this object MUST return a PathErr
message with the "Unknown Object Class" error code (see [RFC2205]).

Also, the following restrictions MUST be applied with respect to the
PPRO usage:

- PPROs MAY only be included in Path messages when signaling
  secondary protecting LSPs (S bit = 1 and P bit = 1) and when the
  LSP Protection Type value is set to 0x02 (Re-routing without
  Extra-Traffic) in the PROTECTION object (see Section 14.).

- PRROs SHOULD be present in the Path message for the pre-
    provisioning of the secondary protecting LSP to enable recovery
    resource sharing between one or more secondary protecting LSPs
    (see Section 15.4).

  - PPROs MUST NOT be used in any other conditions. In particular, if
    a PPRO is received when the S bit is set to 0 in the PROTECTION
    object, the receiving node MUST return a PathErr message with the
    new error code/sub-code "Routing Problem/PRIMARY PATH_ROUTE object
    not applicable".

  - Crossed exchanges of PPROs over primary LSPs are forbidden (i.e.
    their usage is restricted to a single set of protected LSPs).

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

  - The PPRO's content MUST NOT include subobjects coming from other
    PPROs. In particular, received PPROs MUST NOT be re-used to
    establish other working or protecting LSPs.

15.4 Processing

   The PPRO enables sharing recovery resources between a given
   secondary protecting LSP and one or more secondary protecting LSPs
   if their corresponding primary working LSPs have mutually
   (link/node/SRLG) disjoint paths. Consider a node N through which n
   secondary protecting LSPs (say P[1],...,P[n]) have already been
   established and protecting n primary working LSPs (say
   P'[1],...,P'[n]). Suppose also that these n secondary working LSPs
   share a given outgoing link resource (say r).

   Now, suppose that node N receives a Path message for an additional
   secondary protecting LSP (say Q, protecting Q'). The PPRO carried by
   this Path messages is processed as follows:

  - N checks whether the primary working LSPs P'[1],...,P'[n]
    associated with the LSPs P[1],...,P[n] respectively have any link,
    node and SLRG in common with the primary working Q' (associated
    with Q) by comparing the stored PPRO subobjects associated with
    P'[1],...,P'[n] with the PPRO subobjects associated with Q'
    received in the Path message.

  - If this is the case, N SHOULD NOT attempt to share the outgoing
    link resource r between P[1],...,P[n] and Q. However, upon local
    policy decision, N MAY allocate another available (shared) link
    other than r for use by Q. If this is not the case (upon the local
    policy decision that no other link is allowed to be allocated for

Q) or if no other link is available for Q, N SHOULD return a
    PathErr message with the new error code/sub-code "Admission
    Control Failure/LSP Admission Failure".

    - Otherwise (if P'[1],...,P'[n] and Q' are fully disjoint), the link
      r selected by N for the LSP Q MAY be exactly the same as the one
      selected for the LSPs P[1],...,P[n]. This, after verifying (also
      from its local policy) that the selected link r can be shared
      between these LSPs. If this is not the case (for instance, the
      sharing ratio has reached its maximum for that link) and upon
      local policy decision no other link is allowed to be allocated for
      Q, N SHOULD return a PathErr message with the error code/sub-code
      "Admission Control Failure/Requested Bandwidth Unavailable" (see
      [RFC2205]). Otherwise (if no other link is available), N SHOULD
      return a PathErr message with the new error code/sub-code
      "Admission Control Failure/LSP Admission Failure".

    Note that the process, through which m out of the n (m =< n)
    secondary protecting LSPs PPROs may be selected on a local basis to

    perform the above comparison and subsequent link selection, is out
    of scope of this document.

16. ASSOCIATION Object

    The ASSOCIATION object is used to associate LSPs with each other. In
    the context of end-to-end LSP recovery, the association MUST only
    identify LSPs that support the same Tunnel ID as well as the same
    tunnel sender address and tunnel end point address. The Association
    Type, Association Source and Association ID fields of the object
    together uniquely identify an association. The object uses an object
    class number of the form 11bbbbbb to ensure compatibility with non-
    supporting nodes.

    The ASSOCIATION object is used to associate LSPs with each other.

16.1 Format

    The IPv4 ASSOCIATION object (Class-Num of form 11bbbbbb with value =
    198, C-Type = 1, suggested values, TBA by IANA) has the format:

```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |             Length            | Class-Num(TBD)|  C-Type (1)   |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       Association Type        |        Association ID         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                  IPv4 Association Source                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The IPv6 ASSOCIATION object (Class-Num of form 11bbbbbb with value =
198, C-Type = 2, suggested values, TBA by IANA) has the format:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|            Length             | Class-Num(TBD)|  C-Type (2)   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|       Association Type        |        Association ID         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
|                  IPv6 Association Source                       |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

   Association Type: 16 bits

      Indicates the type of association being identified. Note that
      this value is considered when determining association. The
      following are values defined in this document.

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

            Value   Type
            -----   ----
               0   Reserved
               1   Recovery (R)

   Association ID: 16 bits

      A value assigned by the LSP head-end. When combined with the
      Association Type and Association Source, this value uniquely
      identifies an association.

   Association Source: 4 or 16 bytes

      An IPv4 or IPv6 address, respectively, that is associated to
      the node that originated the association.

16.2. Processing

In the end-to-end LSP recovery context, the ASSOCIATION object is
used to associate a recovery LSP with the LSP(s) it is protecting or
a protected LSP(s) with its recovery LSP. The object is carried in
Path messages. More than one object MAY be carried in a single Path
message.

Transit nodes MUST transmit, without modification, any received
ASSOCIATION object in the corresponding outgoing Path message.

An ASSOCIATION object with an Association Type set to the value
"Recovery" is used to identify an LSP Recovery related association.
Any node associating a recovery LSP MUST insert an ASSOCIATION
object with the following setting:
- the Association Type MUST be set to the value "Recovery" in the
  Path message of the recovery LSP
- the (IPv4/IPv6) Association Source MUST be set to the tunnel
  sender address of the LSP being protected
- the Association ID MUST be set to the LSP ID of the LSP being
  protected by this LSP or the LSP protecting this LSP. If unknown,
  this value is set to its own signaled LSP_ID value (default).
  Also, the value of the Association ID MAY change during the
  lifetime of the LSP.

Terminating nodes use received ASSOCIATION object(s) with the
Association Type set to the value "Recovery" to associate a recovery
LSP with its matching working LSP. This information is used to bind
the appropriate working and recovery LSPs together. Such nodes MUST
ensure that the received Path messages including ASSOCIATION
object(s) are processed with the appropriate PROTECTION object
settings, if present (see Section 14 for PROTECTION object
processing). Otherwise, this node MUST return a PathErr message with
the new error code/sub-code "LSP Admission Failure/Bad Association
Type". Similarly, terminating nodes receiving a Path message with a

PROTECTION object requiring association between working and recovery
LSPs MUST include an ASSOCIATION object. Otherwise, such nodes MUST
return a PathErr message with the new error code/sub-code "Routing
Problem/PROTECTION object not Applicable".

17. Updated RSVP Message Formats

This section presents the RSVP message related formats as modified
by this document. Unmodified RSVP message formats are not listed.

The format of a Path message is as follows:

```
<Path Message> ::= <Common Header> [ <INTEGRITY> ]
                   [ [<MESSAGE_ID_ACK> | <MESSAGE_ID_NACK>] ... ]
                   [ <MESSAGE_ID> ]
                   <SESSION> <RSVP_HOP>
                   <TIME_VALUES>
                   [ <EXPLICIT_ROUTE> ]
                   <LABEL_REQUEST>
                   [ <PROTECTION> ]
                   [ <LABEL_SET> ... ]
                   [ <SESSION_ATTRIBUTE> ]
                   [ <NOTIFY_REQUEST> ... ]
                   [ <ADMIN_STATUS> ]
                   [ <ASSOCIATION> ... ]
                   [ <PRIMARY_PATH_ROUTE> ... ]
                   [ <POLICY_DATA> ... ]
                   <sender descriptor>
```

The format of the <sender descriptor> for unidirectional and
bidirectional LSPs is not modified by the present document.

The format of a Resv message is as follows:

```
<Resv Message> ::= <Common Header> [ <INTEGRITY> ]
                   [ [<MESSAGE_ID_ACK> | <MESSAGE_ID_NACK>] ... ]
                   [ <MESSAGE_ID> ]
                   <SESSION> <RSVP_HOP>
                   <TIME_VALUES>
                   [ <RESV_CONFIRM> ]  [ <SCOPE> ]
                   [ <PROTECTION> ]
                   [ <NOTIFY_REQUEST> ]
                   [ <ADMIN_STATUS> ]
                   [ <POLICY_DATA> ... ]
                   <STYLE> <flow descriptor list>
```

   <flow descriptor list> is not modified by this document.

18. Security Considerations

J.P.Lang et al.   Expires February 2007   36

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

The security threats identified in [RFC4426] may be experienced due
to the exchange of RSVP messages and information as detailed in this
document. The following security mechanisms apply.

RSVP signaling MUST be able to provide authentication and integrity. Authentication is required to ensure that the signaling messages are originating from the right place and have not been modified in transit.

For this purpose, [RFC2747] provides the required RSVP message authentication and integrity for hop-by-hop RSVP message exchanges. For non hop-by-hop RSVP message exchanges the standard IPSEC based integrity and authentication can be used as explained in [RFC3473].

Moreover, this document makes use of the Notify message exchange. This precludes RSVP's hop-by-hop integrity and authentication model. In the case, when the same level of security provided by [RFC2747] is desired, the standard IPsec based integrity and authentication can be used as explained in [RFC3473].

To prevent from the consequences of poorly applied protection and increased risk of misconnection, in particular, when Extra Traffic is involved, that would deliver the wrong traffic to wrong destination, specific mechanisms have been put in place as described in Section 7.2, 8.3 and 10.

19. IANA Considerations

IANA assigns values to RSVP protocol parameters. Within the current document a PROTECTION object (new C-Type), a PRIMARY PATH ROUTE object, and an ASSOCIATION object are defined. In addition, new Error code/sub-code values are defined in this document. Finally, registration of the ADMIN_STATUS object bits is requested.

Two RSVP Class Numbers (Class-Num) and three Class Types (C-Types) values have to be defined by IANA in registry:

http://www.iana.org/assignments/rsvp-parameters

1) PROTECTION object (defined in Section 14.1)

o PROTECTION object: Class-Num = 37

   - Type 2: C-Type = 2 (suggested)

2) PRIMARY PATH ROUTE object (defined in Section 15.1)

o PRIMARY PATH ROUTE object: Class-Num = TBA (of form 0bbbbbbb),

   - Primary Path Route: C-Type = 1 (suggested)

3) ASSOCIATION object (defined in Section 16.1)

   o ASSOCIATION object: Class-Num = TBA (of form 11bbbbbb, value 198
     is suggested)

     - IPv4 Association: C-Type = 1 (suggested)
     - IPv6 Association: C-Type = 2 (suggested)

   o Association Type

   The following values defined for the Association Type (16 bits)
   field of the ASSOCIATION object.

            Value   Type
            -----   ----
              0   Reserved
              1   Recovery (R)

   Assignment of values (from 2 to 65535) by IANA are subject to IETF
   expert review process i.e. IETF Standards Track RFC Action.

   4) Error Code/Sub-code values

   The following Error code/sub-code values are defined in this
   document:

   Error Code = 01: "Admission Control Failure" (see [RFC2205])

   o "Admission Control Failure/LSP Admission Failure"
                                              (suggested value = 4)
   o "Admission Control Failure/Bad Association Type"
                                              (suggested value = 5)

   Error Code = 02: "Policy Control Failure" (see [RFC2205])

   o "Policy Control failure/Hard Pre-empted" (suggested value = 20)

   Error Code = 24: "Routing Problem" (see [RFC3209])

   o "Routing Problem/Unsupported LSP Protection"
                                              (suggested value = 17)
   o "Routing Problem/PROTECTION object not applicable"
                                              (suggested value = 18)
   o "Routing Problem/Bad PRIMARY PATH_ROUTE object"
                                              (suggested value = 19)
   o "Routing Problem/PRIMARY PATH_ROUTE object not applicable"
                                              (suggested value = 20)

   Error Code = 25: "Notify Error" (see [RFC3209])

o "Notify Error/LSP Failure"   (suggested value = 6)
o "Notify Error/LSP Recovered"   (suggested value = 7)
o "Notify Error/LSP Locally Failed"   (suggested value = 8)

5) Registration of the ADMIN_STATUS object bits

The ADMIN_STATUS object (Class-Num = 196, C-Type = 1) is defined in [RFC3473].

IANA is also requested to track the ADMIN_STATUS bits extended by this document. For this purpose, the following new registry entries are requested in the registry entry:

http://www.iana.org/assignments/gmpls-sig-parameters

- ADMIN_STATUS bits:

    Name: ADMIN_STATUS bits
    Format: 32-bit vector of bits
    Position:
        [0]   Reflect (R) bit defined in [RFC3471]
        [1..25]   To be assigned by IANA via IETF Standards
                    Track RFC Action.
        [26]   Lockout (L) bit is defined in Section 13
        [27]   Inhibit alarm communication (I) in [ALARM]
        [28]   Call control (C) bit is defined in [GMPLS-
                    CALL]
        [29]   Testing (T) bit is defined in [RFC3471]
        [30]   Administratively down (A) bit is defined in
                    [RFC3471]
        [31]   Deletion in progress (D) bit is defined in
                    [RFC3471]

20. Acknowledgments

The authors would like to thank John Drake for its active collaboration, Adrian Farrel for his contribution to this document (in particular, to the Section 10 and 11) and his thorough review of the document, Bart Rousseau (for editorial review), Dominique Verchere, and Stefaan De_Cnodder. Thanks also to Ichiro Inoue for his valuable comments.

The authors would like also to thank Lou Berger for the time and effort he spent together with the design team, in contributing to the present document.

21. References

21.1 Normative References

   [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.

J.P.Lang et al.   Expires February 2007   39

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

   [RFC2205]   R.Braden (Editor), "Resource ReserVation Protocol --
                Version 1 Functional Specification", RFC 2205,
                September 1997.

   [RFC2747]   F.Baker et al., "RSVP Cryptographic Authentication",
                RFC 2747, October 2000.

   [RFC2961]   L.Berger et al., "RSVP Refresh Overhead Reduction
                Extensions," RFC 2961, April 2001.

   [RFC3209]   D.Awduche et al., "RSVP-TE: Extensions to RSVP for
                LSP Tunnels," RFC 3209, December 2001.

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

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

   [RFC3477]   K.Kompella, and Y.Rekhter, "Signaling Unnumbered Links
                in Resource Reservation Protocol - Traffic Engineering
                (RSVP-TE)," RFC 3477, January 2003.

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

   [RFC4202]   K.Kompella (Editor), " Routing Extensions in Support of
                Generalized Multi-Protocol Label Switching (GMPLS),"
                RFC 4202, October 2005.

    [RFC4204]   J.Lang (Editor), "Link Management Protocol (LMP)," RFC
                 4204, October 2005.

    [RFC4426]   J.P.Lang, B.Rajagopalan, and D.Papadimitriou (Editors),
                 "Generalized MPLS Recovery Functional Specification,"
                 RFC 4426, March 2006.

    [SEGREC]   L.Berger et al., "GMPLS Based Segment Recovery,"
                 Internet Draft, Work in progress, draft-ietf-ccamp-
                 gmpls-segment-recovery-03.txt, October 2006.

21.2 Informative References

    [ALARM]   L.Berger (Editor), "GMPLS - Communication of Alarm
                 Information", Internet draft, Work in progress, draft-
                 ietf-ccamp-gmpls-alarm-spec-06.txt, September 2006.

    [CRANK]   A.Farrel (Editor), "Crankback Signaling Extensions for
                 MPLS and GMPLS Signaling", Internet Draft, Work in
                 progress, draft-ietf-ccamp-crankback-05.txt, May 2005.

    [GMPLS-CALL] D.Papadimitriou and A.Farrel (Editors), "Generalized
                 MPLS (GMPLS) RSVP-TE Signaling Extensions in support of
                 Calls", Internet draft, Work in progress, draft-ietf-
                 ccamp-gmpls-rsvp-te-call-01.txt, August 2006.

    [RFC4090]   P.Pan (Editor), "Fast Reroute Extensions to RSVP-TE for
                 LSP Tunnels," RFC 4090, May 2005.

    [RFC4427]   E.Mannie and D.Papadimitriou (Editors), "Recovery
                 (Protection and Restoration) Terminology for GMPLS,"
                 RFC 4427, March 2006.

    [XRO]   C.Y.Lee et al. "Exclude Routes - Extension to RSVP-TE,"
                 Internet Draft, Work in progress, draft-ietf-ccamp-
                 rsvp-te-exclude-route-05.txt, August 2005.

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

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

22. Editor's Addresses

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

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

        Dimitri Papadimitriou
        Alcatel
        Copernicuslaan 50
        B-2018, Antwerpen, Belgium
        EMail: dimitri.papadimitriou@alcatel.be

23. 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:

draft-ietf-ccamp-gmpls-recovery-e2e-signaling-04.txt   October 2006

        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 (Perceval)
        Rue Tenbosch, 9

1000 Brussels, Belgium
Phone: +32-2-6409194
EMail: eric.mannie@perceval.net

Dimitri Papadimitriou (Alcatel)
Copernicuslaan 50
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

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

This document and the information contained herein are provided on
an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE
REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE
INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR
IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF
THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED
WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Full Copyright Statement

Acknowledgment