RSVP-TE Extensions in support of End-to-End GMPLS-based Recovery

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt

Status of this Memo

Abstract

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

1. **Contributors**

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

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

   Sudheer Dharanikota (Consult)
   E-mail: sudheer@ieee.org

   Jonathan Lang (Rincon Networks)
   E-mail: jplang@ieee.org

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

   Eric Mannie (Consult)
   Email: eric_mannie@hotmail.com

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

   Bala Rajagopalan (Tellium)
   2 Crescent Place - P.O. Box 901
   Oceanport, NJ 07757-0901, USA
   E-mail: braja@tellium.com

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

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-based 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 [TERM]).

   A functional description of GMPLS-based recovery is provided in
   [FUNCT] and should be considered as a companion document to this
   memo which describes the protocol specific procedures for GMPLS
   RSVP-TE (Resource ReSerVation Protocol - Traffic Engineering)
   signaling (see [RFC-3473]) to support end-to-end recovery of an
   entire LSP from the head-end to the tail-end. The present memo
   addresses four types of end-to-end LSP recovery: 1+1 unidirectional/
   1+1 bi-directional protection, LSP protection with extra-traffic
   (including 1:1 protection with extra-traffic), pre-planned LSP re-
   routing without extra-traffic (including shared mesh) and full LSP
   re-routing.

   The simplest notion of end-to-end LSP protection is the 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 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.

Pre-planned LSP restoration or re-routing (without extra-traffic) relies on the establishment between the same end points 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 LSPs are pre-reserved and 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 activated, it can not carry any extra-traffic. Therefore, this mechanism protects against working LSP failure(s) but requires activation of the protecting LSP after 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. Similarly, the recovery resources are pre-reserved and 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.

Last, full LSP restoration or re-routing, on the other hand, switches normal traffic to an alternate LSP fully established after failure occurrence. The new alternate route is selected at the LSP head-end, it may reuse intermediate node's resources of the failed LSP and may include additional intermediate nodes and/or links.

Note that crankback signaling and intermediate LSP recovery are further detailed in dedicated companion documents.

3. **Conventions used in this document:**

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

In addition, the reader is assumed to be familiar with the terminology used in [GMPLS-ARCH], [RFC-3471], [RFC-3473] and referenced as well as [TERM] and [FUNCT].

4. **Identification**

LSP tunnels are identified by a combination of the SESSION and SENDER_TEMPLATE objects (see also [RFC-3209]). 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.

Extended Tunnel ID

> A 32-bit (or 16-byte) identifier used in the SESSION that

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

> 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 LSP tunnels that belong to the same session.

4.2 **Recovery Identification**

This is done using the following fields in the PROTECTION object (see also Section 14).

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 resource allocation and cross-connection have been
  committed). Both working and protecting LSPs can be primary LSPs.
  A secondary LSP is a control plane provisioned only LSP for which
  resource allocation MAY have been done but for which no cross-
  connection has been performed. Only protecting LSPs can be
  secondary LSPs.

- 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 to 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

[draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt](#)   May 2003

primary LSPs (i.e. fully established) whilst primary LSPs can be
either working or protecting LSPs.

[4.2.2](#) **LSP Recovery**

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

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

- (Pre-planned) LSP Re-routing without Extra-traffic: set if the
  protecting LSPs are secondary LSPs that allows for sharing of the
  recovery resources between one or more than one <sender;receiver>
  pair. When secondary LSPs resources are not dedicated to a single
  <sender;receiver> pair, one refers to shared mesh recovery.

- LSP Protection with Extra-traffic: set if the protecting LSPs are
     dedicated primary LSPs that do allow for extra-traffic transport
     and thus precluding any sharing of the recovery resources between
     more than one <sender;receiver> pair. This type includes 1:1 path
     protection with extra-traffic.

   - Dedicated LSP Protection: set if the protecting LSPs do 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). Note also that
     this document makes a distinction between 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 are end-to-
   end significant 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 service. The net result of this is
   that a single bit (the S bit alone) does not allow differentiating
   whether resource allocation should be performed *with respect to*
   the status of the LSP within the protected entity. The introduction
   of the P bit solves unambiguously this problem. 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 transparent delivery of
   protected services since any intermediate node is not required to

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

   know the semantic associated with the incoming LSP Protection Type
   value.

4.2.3 LSP Association

   When used for the working LSP signaling, the Associated LSP ID
   identifies the protecting LSP. When used for the protecting LSP
   signaling, this field identifies the LSP protected by this 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
```
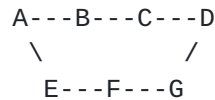
The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint,
ignoring the ingress/egress nodes A and D. A 1+1 protected path is
established from A to D over [A,B,C,D] and [A,E,F,G,D] and traffic
is transmitted simultaneously over both component paths (i.e. LSPs).

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. Note also that both LSPs are instantiated and
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 the present case. Also, for the protected
LSP under failure condition, the Path_State_Remove Flag of the
ERROR_SPEC object (see [RFC-3473]) SHOULD NOT be set upon PathErr
message generation.

Note: one should assume that both paths are SRLG disjoint otherwise,
a failure would impact both working and protecting LSPs.

5.1. Identifiers

Since both LSPs correspond to the same session, 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 is included in the Path message. This object
carries the desired end-to-end LSP Protection Type (in this case,
"1+1 Unidirectional") as well as the LSP ID of the associated LSP

referred to as the Associated LSP ID. This LSP Protection Type value
is applicable to both uni- and 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 the Associated LSP ID to the protecting
LSP_ID. The protecting LSP is signaled by setting in this object the
S bit to 0, the P bit to 1 and the Associated LSP ID to the
associated protected LSP_ID.

After protection switching completes, to keep track of the LSP from
which the signal is taken, the former protecting LSP SHOULD be
signaled as the working LSP from the head-end node (upon reception
of the PathErr message). For the same reason, the former working LSP
SHOULD be signaled as the protecting LSP with the A bit set in the
ADMIN_STATUS object (see [RFC-3473]).

6. 1+1 Bi-directional Protection

1+1 bi-directional protection is another scheme that provides end-
to-end LSP protection.

Consider the following network topology:

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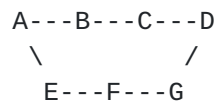
The LSPs [A,B,C,D] and [A,E,F,G,D] are node and link disjoint,
ignoring the ingress/egress nodes A and D. A bi-directional LSP is
established from A to D over each path and traffic is transmitted
simultaneously over both LSPs. In this scheme, both end-points must
receive traffic over the same LSP. When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

Note: one should assume that both paths are SRLG disjoint otherwise
a failure would impact both working and protecting LSPs.

6.1. Identifiers

Since both LSPs correspond to the same session, the SESSION object
MUST be the same in both LSPs. The LSP ID, however, MUST be
different to distinguish between the two LSPs.

A new PROTECTION object is included in the Path message. This object carries the desired end-to-end LSP Protection Type (in this case, "1+1 Bi-directional") as well as the LSP ID of the associated LSP referred to as Associated LSP ID. 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 the Associated LSP ID to the protecting LSP_ID. The protecting LSP is signaled by setting in this object the S bit to 0, the P bit to 1 and the Associated LSP ID to the associated protected LSP_ID.

6.2. End-to-End Switchover Request/Response

To co-ordinate the switchover between endpoints, an end-to-end switchover request is needed since a failure affecting one the LSPs results in both endpoints switching to the LSP (or equivalently the traffic) in their respective direction. This is done using the Notify message with a new Error Code indicating "Working Path Failure; Switchover Request". The Notify Ack message MUST be sent to confirm the reception of the Notify message.

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 [RFC-3473]), it MUST begin receiving on the protection LSP and send a Notify message reliably to the other end-node (D or A, respectively). This message MAY indicate the identity of the failed working link and other relevant information using the IF_ID ERROR_SPEC (see [RFC-3473]).

   Note: in this case, the IF_ID ERROR_SPEC replaces the ERROR_SPEC in the Notify message, otherwise the corresponding (data plane) information is to be received in the PathErr/ResvErr message.

2. Upon receipt of the switchover message, the end-node (D or A, respectively) MUST begin receiving from the protection LSP and send a (Notify) Ack message to the other end-node (A or D, respectively) using reliable message delivery (see [RFC-2961]).

Since the intermediate nodes (B,C,E,F and G) are assumed to be GMPLS 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 working LSP under failure, the
Path_State_Remove Flag of the ERROR_SPEC object (see [RFC-3473])
SHOULD NOT be set upon PathErr message generation. After protection
switching completes (step 2), to keep track of the LSP from which
the signal is taken, the former protecting LSP SHOULD be signaled as
the working LSP. For the same reason, the former working LSP SHOULD
be signaled as the protecting LSP with the A bit set in the
ADMIN_STATUS object (see [RFC-3473]).

Note: 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).

7. 1:1 Protection with Extra-Traffic

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

An important feature of GMPLS signaling is that it allows pre-
provisioning of protecting LSPs to protect working LSPs. This is
done by indicating in the Path message (in the newly defined
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 both signaled as primary LSPs; they are fully
instantiated during the provisioning phase.

Although the resources for the protecting LSPs are pre-allocated
lower priority traffic may use these resources (i.e. the protecting
LSP are capable to carry extra-traffic) with the caveat that the
lower priority traffic will be preempted if the working LSP fails.
If lower priority traffic is using resources along the protecting
LSPs, the end-nodes may need to be notified of the failure in order
to complete the switchover.

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 upstream node (upstream in terms of the
direction an RSVP Path message traverses) SHOULD send an RSVP Notify
message to the LSP head-end, and the downstream node SHOULD send an
RSVP Notify message to the LSP tail-end. 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). Note that if the working and the protecting LSP are

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

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
```
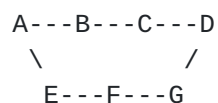

The working LSP [A,B,C,D] could be protected by the protecting LSP
[A,E,F,G,D]. Both LSPs are instantiated (resources are allocated for
both working and protecting LSPs) and no resource sharing can be
done along the protection path since the primary protecting LSP can
carry extra-traffic.

Note: one should assume that both paths are SRLG disjoint otherwise
a failure would impact both working and protecting LSPs.

7.1 Identifiers

Since both LSPs correspond to the same session, the SESSION object
MUST be the same in 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 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, "1:1 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 this object the S bit to
0, the P bit to 0 and the Associated LSP ID to the protecting
LSP_ID. The protecting LSP is signaled by setting in this object the
S bit to 0, the P bit to 1 and the Associated LSP ID to the

associated protected LSP_ID.

7.2 End-to-End Switchover Request/Response

To co-ordinate the switchover between endpoints, an end-to-end
switchover request is needed such that the affected LSP(s) are 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-point nodes
(A or D) or simply end-nodes.

This operation may be done using Notify message exchange with a new
Error Code indicating "Working Path Failure; Switchover Request".
The Notify Ack message MUST be sent to confirm the reception of the
Notify message.

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
   [RFC-3473]), it disconnects the extra-traffic from the
   protecting LSP and send a Notify message reliably to the
   other end-node (D or A, respectively). This message MAY
   indicate the identity of the failed working link and other
   relevant information using the IF_ID ERROR_SPEC (see [RFC-
   3473]).

   Note: in this case, the IF_ID ERROR_SPEC replaces the
   ERROR_SPEC in the Notify message, otherwise the corresponding
   information is to be received in the PathErr/ResvErr message

2. Upon receipt of the switchover (i.e. 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 and send a
   (Notify) Ack message to the other end-node (A or D,
   respectively) using reliable message delivery (see [RFC
   2961]). Also, the Notify message generated by the end-node
   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 (Notify) Ack message, the end-node (A or
      D, respectively) MUST begin receiving normal traffic from
      the protecting LSP.

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

   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), the formerly working
   LSP SHOULD be signaled with the A bit set in the ADMIN_STATUS object
   (see [RFC-3473]).

8. 1:1 Re-routing without Extra-Traffic

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

   End-to-end LSP 1:1 re-routing without Extra-Traffic relies on the
   establishment between the same endpoints of a working LSP and a
   protecting LSP that is link/node/SRLG disjoint from the working one.
   However, in this case the protecting LSP is not instantiated, thus,
   it can not carry any extra-traffic. Therefore, this mechanism
   protects against working LSP failure(s) but requires instantiation
   of the protecting LSP after failure occurrence.

   Signalling is performed by indicating in the Path message (in the
   newly defined 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 pre-allocated (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 [RFC-3473]) 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.

Consider the following topology:


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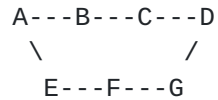

The working LSP [A,B,C,D] could be protected by the protecting LSP
[A,E,F,G,D]. Only the protected LSP is instantiated (resources are
only allocated for the working LSP) therefore, the protecting LSP
can not carry any extra-traffic. When a failure is detected on the
working LSP (say at B), the error is propagated and/or notified to
the ingress node (A), which activates the secondary protecting LSP
instantiated during the 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.

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

8.1 Identifiers

Since both LSPs correspond to the same session, the SESSION object
MUST be the same in both LSPs. The LSP ID, however, MUST be
different to distinguish between the two LSPs, here the protected
LSP carrying working traffic and the protecting LSP that can not
carry extra-traffic.

A new PROTECTION object is used to setup the two LSPs. This object
carries the desired end-to-end LSP Protection Type in this case,
"1:1 Re-routing without Extra-Traffic") as well as the LSP ID of the
associated LSP. This LSP Protection Type value is applicable to both
uni- and bi-directional LSPs.

8.2 Signaling Primary LSPs

The PROTECTION object is included in the Path message during
signaling of the primary working LSP, with the end-to-end LSP
Protection Type set to "1:1 Re-routing without Extra-Traffic". The
primary working LSP is signaled by setting in this object the S bit
to 0, the P bit to 0 and the Associated LSP ID to the protecting
LSP_ID.

8.3 Signaling Secondary LSPs

Secondary LSPs are signaled using the S bit of the new PROTECTION
object that is carried. If set, 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 Path refresh message
with the S bit set to 0 in the PROTECTION object. At this point, the
link and node resources must to be allocated for the LSP that
becomes a primary working LSP (ready to carry normal traffic).

Two cases have to be covered here (see also [GMPLS-ARCH]) since
secondary protecting LSPs can be setup with resource reservation but
with or without label pre-selection (both allowing sharing of the
recovery resources). In the former case (defined as the default),
secondary LSP signaling does not necessitate any specific procedure
compared to the one defined in [RFC-3473]. However, in the latter
case, label (and thus resource) re-allocation MAY occur during the
secondary LSP activation. This means that during the activation
phase, labels MAY be re-assigned (with higher precedence over label
assignment, see also [RFC-3471]).

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 [FUNCT]. Shared-mesh restoration can be seen as
particular case of pre-planned LSP re-routing that reduces the
recovery resource requirements by allowing multiple working LSPs to
share common link and node resources. Here also, the recovery
resources for the protecting LSPs are pre-reserved during the
provisioning phase, but explicit (signaling) action is required to
activate (i.e. commit resource allocation at the data plane) a
specific protecting LSP instantiated during the provisioning phase.

This requires restoration signaling along the protecting path.

Consider the following topology:

```
A---B---C---D
 \         /
  E---F---G
 /         \
H---I---J---K
```
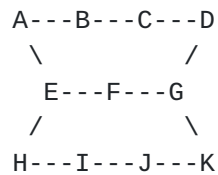
The working LSPs [A,B,C,D] and [H,I,J,K] could be protected by
[A,E,F,G,D] and [H,E,F,G,K], respectively. In order to achieve
resource merging during the signaling of these recovery LSPs (i.e.
resource sharing), the LSPs must have the same Session Ids, but the
Session Id includes the target (egress) IP address. 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 Type
setting of the secondary LSPs is for protection (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 can be shared.

When a failure is detected on one of the working LSPs (say at B),
the error is propagated and/or notified to the ingress node (A),
which activates the 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 [RFC-3473]). 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.

10.1. Identifiers

   Since both LSPs (i.e. the primary working and the secondary
   protecting LSPs) correspond to the same session, the SESSION object
   MUST be the same for both LSPs. The LSP ID, however, MUST be
   different to distinguish between the two LSPs.

10.2 Signaling Primary LSPs

   A new PROTECTION object is included in the Path message during
   signaling of the primary working LSP. The PROTECTION object carries
   the desired end-to-end LSP Protection Type (in this case, "1:1 Re-
   routing without Extra-Traffic") as well as the LSP ID of the
   associated protecting LSP. This LSP Protection Type value is
   applicable to both uni- and bi-directional LSPs.

   Primary working LSPs are signaled by setting the both S bit and the
   P bit of the PROTECTION object to 0.

10.3 Signaling Secondary LSPs

   The new PROTECTION object carried in the Path message includes the
   desired end-to-end LSP Protection Type (in this case, "1:1 Re-
   routing without Extra-Traffic") as well as the LSP ID of the
   associated primary protected LSP, which MUST be known before
   signaling of the secondary LSP. This LSP Protection Type value is
   applicable to both uni- and bi-directional LSPs.

   Secondary LSPs are signaled by setting in this object the S bit to 1
   and the P bit to 1. Moreover, the Path message used to instantiate
   the secondary LSP MUST include at least one PRIMARY PATH ROUTE
   object (see Section 15) that enables distinguishing shared mesh
   restoration at each intermediate node along the secondary path.

   Secondary LSPs are signaled using the S bit of the new PROTECTION
   object that is carried in the Path message. If set, 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 Path
refresh message with the S bit set to 0 in the PROTECTION object. At
this point, the link and node resources must to be allocated for the
LSP that becomes a primary working LSP (ready to carry normal
traffic).

Two cases have to be covered here (see also [GMPLS-ARCH]) since the
secondary LSP can be setup with resource reservation but with or
without label pre-selection (both allowing sharing of the recovery
resources). In the former case (defined as the default), secondary
LSP signaling does not necessitate any specific procedure compared
to the one defined in [RFC-3473]. 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 label assignment,
see also [RFC-3471]).

11. (Full) LSP Re-routing

LSP re-routing, on the other hand, switches normal traffic to an
alternate LSP that is fully established 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 filed 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.

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 failure or received either a Notify message
and/or a PathErr message indicating that a failure has occurred. 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 object and
the Shared-Explicit (SE) reservation style (see [RFC-3209]). 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

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

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 and LSP Id are all used to uniquely identify both the old and new LSPs. The new (alternate) LSP is setup before the old LSP is torn down using Shared-Explicit (SE) reservation style. This ensures that the new LSP is established without double counting resource requirements along common segments.

11.2 Signalling 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 this object the S bit to 0, the P bit to 0 and the Associated LSP ID to 0. Any specific action to be taken during the provisioning phase is up 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 MUST NOT be re-routed at the head-end node after failure occurrence. The Associated LSP_ID value MUST be set to 0.

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 remains allocated to the LSP that was originally routed over it 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" and "1:1 with Extra-traffic" protection, reversion to the recovered LSP simply occurs by clearing its A bit in the ADMIN_STATUS object and applying the reverse 1-phase APS

switchover request/response (or 2-phase APS) described in Section
6.2 (or Section 7.2, respectively).

For "Re-routing without Extra-traffic" reversion implies that the
formerly working LSP has not been torn down by the head-end upon
PathErr message reception (i.e. the head-end node kept refreshing
the working LSP under failure condition by setting A bit in the
ADMIN STATUS object). This ensures that the same resources are
retrieved after reversion switching. Re-activation is performed by
clearing the A bit for the recovered working primary LSP and then
set the S bit to 1 in the PROTECTION object sent over the protecting
path.

13. External Commands

This section specifies the control plane behavior when using several
external commands (see [TERM]), typically issued by an operator
through the Network Management System (NMS)/Element Management
System (EMS), which can be used to influence or command the recovery
operations. Other specific commands may complete the below list.

A. Lockout of recovery LSP:

The Administratively Down bit (A bit) of the ADMIN_STATUS object is
used following the rules defined in Section 8 of [RFC-3471] and
Section 7 of [RFC-3473]. The A bit must be set together with the
Reflect (R) bit set in the ADMIN_STATUS object. Its usage forces the
recovery LSP to be temporarily unavailable to transport traffic
(either normal or extra traffic). Unlock is performed by clearing
the A bit.

B. Lockout of normal traffic:

The A bit usage forces the recovery LSP to be temporarily
unavailable to transport normal traffic. The A bit must be set
together with the Reflect (R) bit set in the ADMIN_STATUS object.
Unlock is performed by clearing the A bit.

C. Forced switch for normal traffic:

Recovery signaling is initiated externally that switches normal
traffic to the recovery LSP following the procedure defined in
Section 7.

D. Manual switch for normal traffic:

Recovery signaling operation is initiated externally that switches
normal traffic to the recovery LSP following the procedure defined
in Section 7. This, unless a fault condition exists on other
LSPs/spans (including the recovery LSP) or an equal or higher
priority switch command is in effect.

E. Manual switch for recovery LSP:

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

14. PROTECTION Object

In this section, we describe the extensions to the PROTECTION object
to broaden its applicability to end-to-end LSP recovery. In addition
to modifications to the format of the PROTECTION object, we extend

its use so that the object can be included in the Notify message to
act a switchover request for 1+1 bi-directional and 1:1 protection.

The format of the PROTECTION Object (Class-Num = 37, C-Type = TBA by
IANA) 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(37) | C-Type (TBA)  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |S|P|N|  Reserved   | LSP Flags |     Reserved      | Link Flags|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |      Associated LSP ID        |          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 Flag is set 0x10, 0x08 or 0x04 and MUST be set to
0 in any other case.

Reserved: 7 bits

This field is reserved. It MUST be set to zero on transmission
and MUST be ignored on receipt. These bits SHOULD be pass
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.

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

        0x00   Unprotected
        0x01   (Full) Re-routing
        0x02   1:1 Re-routing without Extra-Traffic
        0x04   1:1 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 pass
through unmodified by transit nodes.

Link Flags: 6 bits

Indicates the desired link protection type (see [RFC-3471]).

Associated LSP ID: 16 bits

Identifies the LSP protected by this LSP or the LSP protecting
this LSP. If unknown, this value is set to 0 (default). Also,
the value of the Associated LSP ID MAY change during the
lifetime of the LSP.

   Reserved: 16 bits

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

Intermediate nodes processing a Path message containing a PRIMARY
PATH ROUTE object (see Section 15) and a PROTECTION object with the
LSP Protection Type "0x02" value set MUST verify that the requested
LSP Protection Type can be satisfied by the outgoing interface. If
it cannot, the node MUST generate a PathErr message, with a "Routing
problem/Unsupported LSP Protection" indication. If due to a resource
unavailability on the outgoing interface, an intermediate node MUST
return a PathErr with the "Routing Problem/LSP Admission Failure"
error code.

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

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 Associated LSP ID field). This object MUST be
if and only if the LSP Protection Type value is set to "0x02". This
memo does not assume 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 LSR that initiates the request.
Therefore, the information included in these objects MAY be used as
policy-based admission control to ensure that secondary protecting
LSPs that are sharing resources have (link/node/SRLG) disjoint paths

for their associated primary LSPs.

15.1. Definition

The primary path route is specified via the PRIMARY_PATH_ROUTE object (PPRO). The Primary Path Route Class Number 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, C-Type = 1

```
  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. The subobjects are identical to those that can constitute an EXPLICIT ROUTE object as defined in [RFC-3209], [RFC-3473] and [RFC-3477].

To signal a secondary protecting LSP, the Path message MUST include at least one or MAY include 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 Applicability

The PRIMARY_PATH_ROUTE object MUST only be used when all GMPLS nodes along the path support the PRIMARY_PATH_ROUTE object and secondary protecting LSPs are 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.

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

- PPROs MUST only be sent over secondary protecting LSPs (S bit = 1

and P bit = 1) and when the LSP Protection Type value is set to
"0x02" in the PROTECTION object (see Section 15.)

   - Crossed exchanges of PPROs over primary LSPs are forbidden (i.e.
     their usage is restricted to a single set of protected LSPs). If a
     PPRO is received with the S bit set to 0 in the PROTECTION object,
     the receiving node MUST return a PathErr with the "Routing
     Problem/PRIMARY PATH_ROUTE object not applicable" error code.

   - 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.

15.3 Subobjects

   The PRIMAY_PATH_ROUTE object is defined as a list of variable-length
   data items called subobjects. PPR subobjects are derived from the
   subobjects of the EXPLICIT ROUTE and/or RECORD ROUTE object of the
   primary working LSP(s). Each PPR subobject has its own length field.
   The length contains the total length of the subobject in bytes,
   including the Type and 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 [RFC 3209])
   - Sub-Type 2: IPv6 Address (see [RFC 3209])
   - Sub-Type 3: Label (see [RFC-3473])
   - Sub-Type 4: Unnumbered Interface (see [RFC-3477])

   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 with the
   "Routing Problem/Bad PRIMARY PATH_ROUTE object" error code.

   Note: SRLG identifier values can be derived from the local IGP-TE
   database using the Type 1, 2 or 4 subobjects listed here above as
   pointers to the corresponding TE Link Id.

16. Application Examples

   This section illustrates the use of the above-defined objects with
   respect to each of the recovery mechanisms considered in this memo.

16.1 1+1 Bi-directional Protection

The protected LSP is signaled with both S bit and P bit set to 0.
The protecting LSP is signaled with the S bit to 0 and P bit set to
1. LSP Flag is set to 0x10 (for both LSP setup). Associated LSP_IDs
point the one to each other.

16.2 1+1 Unidirectional Protection

The protected LSP is signaled with both S bit and P bit set to 0.
The protecting LSP is signaled with S bit set to 0 and P bit set to
1. LSP Flag is set to 0x08 (for both LSP setup). Associated LSP_IDs
point the one to each other.

16.3 1:1 Protection with Extra-Traffic (Path and bandwidth protection)

The protected LSP is signaled with both S bit and P bit set to 0.
LSP Flag is set to 0x04 (during LSP setup). Associated LSP ID points
to the protecting LSP ID.

The protecting LSP (carrying extra-traffic) is signaled with S bit
set to 0 and P bits set to 1. LSP Flag is set to 0x04 (during LSP
setup). Associated LSP ID points to the protected LSP ID.

16.4 1:1 Re-Routing without Extra-Traffic (Path protection only)

The protected LSP is signaled with both S bit and P bit set to 0.
LSP Flag is set to 0x02 (during LSP setup). Associated LSP ID points
to the protecting LSP ID.

The protecting LSP is signaled with both S bit and P bit set to 1.
LSP Flag is set to 0x02 (during LSP setup). Associated LSP ID points
to the protected LSP ID.

16.5 Shared Mesh

Each protected LSP is setup with both S and P bits set to 0. LSP
Flag is set to 0x02 (during LSP setup). Each Associated LSP ID
points to the single protecting LSP ID.

The single protecting LSP is setup with S bit set to 1 and P bits
set to 1. LSP Flag is set to 0x02 (during LSP setup). Associated LSP
ID MUST be set either to the protected LSP (single protected LSP) or
to 0 (multiple Protected LSPs). In addition, the protecting LSP path
message MUST carry at least PPRO object, typically one for each
protected LSP.

16.6 Full Re-routing

Each re-routable LSP is setup with both S and P bits set to 0. LSP
Flag is set to 0x01 (during LSP setup). Associated LSP ID MUST be

set to 0.

17. **Security Considerations**

   This document does not introduce or imply any specific security
   consideration.

18. **Acknowledgments**

   The authors would like to thank John Drake for his active
   collaboration and Adrian Farrel for his contribution to this
   document (in particular to the Section 11). Many thanks also to Bart
   Rousseau (for its editorial revision) and Stefaan De_Cnodder.

19. **IANA Considerations**

   IANA assigns values to RSVP protocol parameters. Within the current
   document a PROTECTION object (new C-Type) and a PRIMARY PATH ROUTE
   object are defined.

   One RSVP Class Number (Class-Num) and two Class Types (C-Types)
   values have to be defined by IANA in registry:

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

   - PROTECTION object: Class-Num = 37, C-Type = 2 (suggested)

   - PRIMARY PATH ROUTE object: Class-Num = 23 (suggested), C-Type =
     1 (suggested)

   - Error codes:

     o "Routing Problem/Unsupported LSP Protection"   (value = TBA)
     o "Routing Problem/LSP Admission Failure"   (value = TBA)
     o "Routing Problem/Bad PRIMARY PATH_ROUTE object" (value = TBA)
     o "Routing Problem/PRIMARY PATH_ROUTE object not applicable"

20. **Intellectual Property Considerations**

   This section is taken from Section 10.4 of [RFC2026].

   The IETF takes no position regarding the validity or scope of any
   intellectual property 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; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in BCP-11. Copies of claims of rights made available for publication 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 implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights, which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

21. References

21.1 Normative References

[FUNCT]   J.P.Lang and B.Rajagopalan (Editors), "Generalized MPLS
              Recovery Functional Specification," Internet Draft,
              Work in Progress, draft-ietf-ccamp-gmpls-recovery-
              functional-00.txt, January 2002.

[GMPLS-ARCH] E.Mannie (Editor), "Generalized MPLS Architecture",
              Internet Draft, Work in progress, draft-ietf-ccamp-
              gmpls-architecture-06.txt, April 2003.

[GMPLS-RTG]   K.Kompella (Editor), "Routing Extensions in Support of
              Generalized MPLS," Internet Draft, Work in Progress,
              draft-ietf-ccamp-gmpls-routing-05.txt, August 2002.

[LMP]   J.Lang (Editor), "Link Management Protocol (LMP) v1.0,"
              Internet Draft, Work in progress, draft-ietf-ccamp-lmp-
              08, March 2003.

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

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

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

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

    [RFC-3471]   L.Berger, (Editor) et al., "Generalized MPLS û
                  Signaling Functional Description," RFC 3471, February
                  2003.

    [RFC-3473]   L.Berger (Editor) et al., "Generalized MPLS
                  Signaling û RSVP-TE Extensions," RFC 3473, February
                  2003.

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

draft-lang-ccamp-gmpls-recovery-e2e-signaling-01.txt   May 2003

    [TERM]   E.Mannie and D.Papadimitriou (Editors), "Recovery
                  (Protection and Restoration) Terminology for GMPLS,"
                  Internet Draft, Work in progress, draft-ietf-ccamp-
                  gmpls-recovery-terminology-02.txt, May 2003.

21.2 Informative References

    [CCAMP-LI]   G.Li et al. "RSVP-TE Extensions for Shared-Mesh
                  Restoration in Transport Networks," Internet Draft,
                  Work in progress, draft-li-shared-mesh-restoration-
                  01.txt, November 2001.

22. Author's Addresses

    Jonathan Lang (Rincon Networks)
    E-mail: jplang@ieee.org

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

Full Copyright Statement

HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF
MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE."