CCAMP Working Group
Internet Draft
Expiration Date: August 2003

CCAMP GMPLS P&R Design Team

J.P. Lang (Editor)
Y. Rekhter (Editor)

February 2003

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

draft-lang-ccamp-gmpls-recovery-e2e-signaling-00.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. Besides the editors, 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

John Drake (Calient)
25 Castilian Drive
Goleta, CA 93117, USA
E-mail: jdrake@calient.net

Jonathan Lang (Consult)
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 MPLS (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. 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 period (see [TERM]).

A functional description of GMPLS-based recovery is provided in [FUNCT] and should be considered a companion document to this document.

This document describes protocol specific procedures for GMPLS (Generalized Multi-Protocol Label Switching) RSVP-TE (Resource ReSerVation Protocol - Traffic Engineering) signaling (see [RFC-3473] to support end-to-end recovery of an entire LSP from the initiator to the terminator. In this memo, we address three types of end-to-end recovery schemes: 1+1 unidirectional protection, 1+1 bi-directional protection, 1:1 protection, and shared mesh restoration.

The simplest notion of end-to-end protection is 1+1 unidirectional protection. In this scheme, a protection (primary) LSP is signaled over a dedicated resource-disjoint alternate path to protect the working (primary) LSP. 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 protection (primary) LSP is signaled over a dedicated resource-disjoint alternate path to protect the working (primary) LSP. 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 a protection LSP.

Shared-mesh restoration reduces the pre-provisioned resource requirements by allowing multiple LSPs to share common link and node resources. In this scheme, the recovery capacity is pre-reserved, but explicit action is required to activate (i.e. commit resource allocation) a specific recovery LSP instantiated during the provisioning phase. This requires restoration signaling along the protection path.

Note that crankback and other intermediate recovery signalling will be addressed in a companion document.

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. LSP 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 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.

5. 1+1 Unidirectional Protection

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

Consider the following network topology:

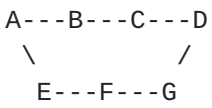

The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A 1+1 protected LSP is established from A to D over [A,B,C,D] and [A,E,F,G,D] and traffic is transmitted simultaneously over both paths (i.e. "LSPs").

When a failure is detected on one path (say at node B), the receiver at D simply selects the traffic from the other LSP. Note that both LSPs are instantiated and no resource sharing can be done along the protection path.

Note: If a failure occurs for instance between link B-C, one should assume that both paths are SRLG disjoint otherwise such a failure

would impact both working and protection LSPs.

5.1. 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 used to setup the two LSPs. 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.

6. 1+1 Bi-directional Protection

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

Consider the following network topology:

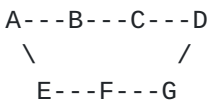

The paths [A,B,C,D] and [A,E,F,G,D] are node and link disjoint, ignoring the ingress/egress nodes A and D. A bi-directional LSP is established from A to D over each path and traffic is transmitted simultaneously over both LSPs. In this scheme, both end-points must receive traffic over the same LSP. When 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. Note that both LSPs are instantiated and no resource sharing can be done along the protection path.

Note: If a failure occurs for instance between link B-C, one should assume that both paths are SRLG disjoint otherwise such a failure would impact both working and protection 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 LSP 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 and referred to as Associated 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 of one the paths results in both endpoints switching to the path (or equivalently the traffic) in their respective direction. This may be done using the Notify message with a new Error Code indicating "Working Path Failure; Switchover Request". The Notify Ack message MUST be sent confirming receipt 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 signalling capable, it has to emphasized that each of them MAY also generate a Notify message directed either to the LSP initiator (upstream direction) or the LSP terminator (downstream direction) or even both. Therefore, it is expected that these LSP terminating

nodes (that also detects the failure of the LSP from the data plane) provides either the right correlation mechanism to avoid repetition of the above procedure or just discard subsequent Notify messages corresponding to the same Session.

Also for 1+1 protected LSP, the Path_State_Remove Flag of the ERROR_SPEC object (see [RFC-3473] for more details) SHOULD NOT be set.

7. 1:1 Dedicated Protection (with Extra Traffic)

The most common notion of 1:1 path protection is to route a node-disjoint primary working LSP and a pre-establish protecting LSP that is link/node/SRLG disjoint from the primary one. This protects against working LSP failure(s).

An important feature of GMPLS signalling is that it allows pre-configuring protecting LSPs to protect working LSPs. This is done by indicating in the Path message (in the newly defined PROTECTION object) that the LSP is of type working and protecting, respectively. Protecting LSPs are used for fast switchover when working LSPs fail. Note also that both working and protecting LSPs are primary LSPs.

Although the resources for the protecting LSPs are pre-allocated, lower priority traffic may use the resources 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 initiator and terminator 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 initiator, and the downstream node SHOULD send an RSVP Notify message to the LSP terminator. Upon receipt of the Notify messages, the initiator and terminator nodes MUST switch the traffic from the working LSP to the pre-configured protecting LSP. Note that if a common initiator-terminator is used for the working and protecting LSPs no further notification is required to indicate

that the working LSPs are no longer protected. Note also that both working and protecting LSPs are working LSPs since fully instantiated during the provisioning phase.
Consider the following topology:

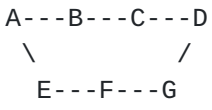

The path [A,B,C,D] could be protected by [A,E,F,G,D]. Both LSPs are instantiated (resources are allocated for both working and protecting LSPs) and no resource sharing can be done along the protection path since the primary protecting LSP can carry extra-traffic.

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 two LSPs, here the protected LSP carrying working traffic and the protection LSP that may 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 with Extra-Traffic") for the working LSP by setting both Protection bit and Secondary bit to 0. The protection LSP is signaled by setting the Protection bit to 1 and the Secondary bit to 0 as well as the LSP ID of the associated working LSP in the PROTECTION object carried with the Path message.

7.2 End-to-End Switchover Request/Response

To co-ordinate the switchover between endpoints, an end-to-end switchover request is needed the affected LSP(s) must be 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 confirming receipt 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]).

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: a 2-phase Automatic Protection Switching (APS) is used in the present context, 3-phase APS (see [FUNCT]) implying a notification message and a switchover request/response messages, are left for further study.

8. End-to-End Bulk Recovery

TBD.

9. Shared Mesh Restoration

An approach to reduce the pre-provisioned resource requirements for recovery 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]. With shared mesh restoration, the capacity for the protection LSPs is pre-reserved and explicit action is required to instantiate the protection LSP.

Consider the following topology:

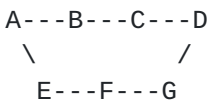

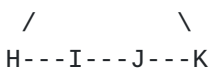

The (working) paths [A,B,C,D] and [H,I,J,K] could be protected by [A,E,F,G,D] and [H,E,F,G,K], respectively. In order to achieve resource merging during the signalling of these recovery LSPs (i.e. resource sharing), the LSPs must have the same Session Ids, but the Session Id includes the target (egress) 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) and acts accordingly. In this case, the recovery 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 primary working path (say at B), the error is propagated to the ingress (A) which instantiates the protection path. At this point, it is important that a failure on the other path (say at J) does not cause the other ingress (H) to send the data down the protection path since the resources are already in use. This can be achieved by node E in two ways. 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. Second, when a failure does occur (say at B) and the protecting LSP is instantiated, E should notify H that the resources for the protecting LSP 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" 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 establishing the secondary LSP
(4) the ability to instantiate a secondary LSP as an active LSP when a failure occurs, and
(5) the ability to instantiate several secondary LSPs as activated LSPs in an efficient manner.

In the following subsections, these features are described in more

detail.

9.1. 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.

9.2 Signaling Primary LSPs

A new PROTECTION object is included in the Path message during signaling of the primary working LSP. These LSPs are signaled by setting the Secondary bit of the PROTECTION object to 0. The PROTECTION object carries the desired end-to-end LSP Protection Type (in this case, "Shared Mesh") as well as the LSP ID of the associated LSP if available and otherwise set to 0.

Note also that in the present context, the protected LSP is considered as working such that the Protection bit of the PROTECTION object is also set to 0.

9.3 Signaling Secondary LSPs

Secondary LSPs are signaled using the Secondary bit of the PROTECTION object that is carried in the Path message. If set, the resources for the secondary LSP should be 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 the secondary LSP. Activation of a secondary LSP is done using a Path refresh message with the "Secondary" bit cleared. At this point, the link and node resources need to be allocated for the LSP.

Moreover, when used for shared mesh recovery purposes, secondary LSPs are signaled using the Protection bit of the PROTECTION object. This object carries the desired end-to-end LSP Protection Type (in this case, "Shared Mesh") as well as the LSP ID of the associated primary LSP, which MUST be known before signaling of the secondary LSP.

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, secondary LSP signalling does not necessitate any specific procedure compared to the one defined in [RFC-3473]. However, in the latter one, 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]).

10. Full LSP Restoration

Full LSP restoration, on the other hand, switches traffic to an alternate route around a failure. The new (alternate) route is selected at the LSP initiator and may reuse intermediate nodes included in the original LSP route; it may also include additional intermediate nodes. For strict-hop routing, TE requirements can be directly applied to the route calculation, and the filed node or link can be avoided. However, if the failure occurred within a

loose-routed hop, the source node may not have enough information to reroute the LSP around the failure.

The alternate route may be calculated on demand (that is, when the failure occurs) or may be pre-calculated and stored for use when the failure is reported. This 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 restoration will be initiated by the node that has isolated the failure or by the node that has received either an RSVP Notify message or an RSVP PathErr message indicating that a failure has occurred. The new resources can be established in a make-before-break fashion, where the new primary LSP is setup before the old primary LSP is torn down. This is done by using the mechanisms of the LSP_Tunnel Session object (see [RFC-3209]) and the Shared-Explicit reservation style. Both the new and old primary LSPs share resources at nodes common to both LSPs. The Tunnel end point addresses, Tunnel Id, Extended Tunnel Id, Tunnel sender address, and LSP Id are all used to uniquely identify both the old and new LSPs; this ensures new resources are established without double counting resource requirements along common segments.

Note that make-before-break is not used to avoid disruption to the data flow (this has already been broken by the failure that is being repaired), but is valuable to retain the resources allocated on the original primary LSP that will be re-used by the new primary LSP.

11. Reversion

TBD.

12. 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/span:

A Lockout bit (L) is defined in the ADMIN_STATUS object that follows the rules defined in Section 8 of [RFC-3471] and Section 7 of [RFC-3473]. Its usage forces the recovery LSP/span to be temporarily unavailable to transport traffic (either normal or extra traffic).

B. Lockout of normal traffic:

The Lockout bit (L) usage results in the normal traffic being temporarily not allowed to be routed over its recovery LSP/span.

C. Freeze:

TBD.

D. Forced switch for normal traffic:

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

E. Manual switch for normal traffic:

Recovery signalling is initiated externally that switches normal traffic to the recovery LSP/span following the procedure defined in Section 7. This, unless a fault condition exists on other LSPs/spans (including the recovery LSP/span).

13. PROTECTION Object

In this section, we describe extensions to the PROTECTION object to extend 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 and 1:1 bi-directional protection. The format of the PROTECTION Object (Class-Num = 37, C-Type = TBA by IANA) is as follows:

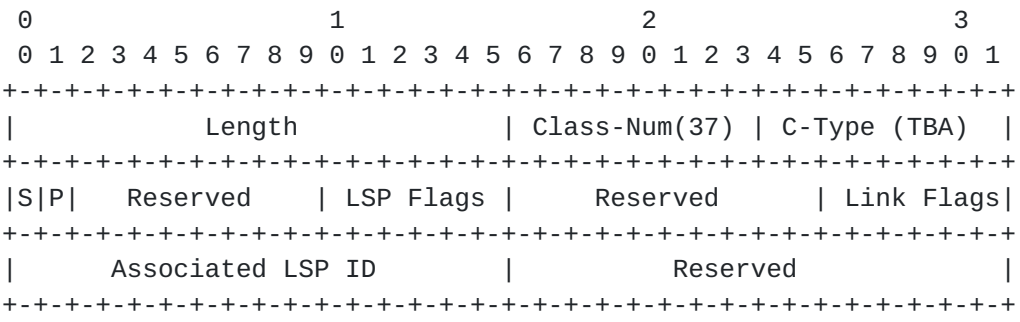

Secondary (S): 1 bit

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 (or recovery) LSP. When set to 0 (default), it indicates that the requested LSP is a working LSP.

Reserved: 8 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) Flags: 6 bits

Indicates desired end-to-end LSP recovery type. A value of 0 implies that LSP recovery type is left unspecified. Only one bit can be set at a time. The following values are defined. All other values are reserved and must be sent as zero and ignored on receipt.

0x00 Unspecified

```
0x01    Extra-Traffic
0x02    Unprotected
0x04    Shared Mesh
0x08    Dedicated 1:1 (with Extra Traffic)
0x10    Dedicated 1+1 Unidirectional
0x20    Dedicated 1+1 Bidirectional
```

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. If unknown, this value is by default set to 0.

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.

14. PRIMARY PATH ROUTE Object

The PRIMARY PATH (Explicit) ROUTE object (PRRO) is defined to inform nodes along the path of a secondary LSP about which resources (link/nodes) are being used by the associated primary LSP (as specified by the Associated LSP ID field). This object MAY also be used to inform nodes along the path of a primary protecting LSP about which resources are being used by the associated primary working LSP.

PRR objects carry the EXPLICIT ROUTE object (see [RFC-3209]) of the LSPs they protect. Therefore, the information included in these objects MAY be used as policy-based admission control to ensure that secondary LSPs that are sharing resources have (link/node/SRLG) disjoint paths for their associated primary LSPs.

14.1. 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

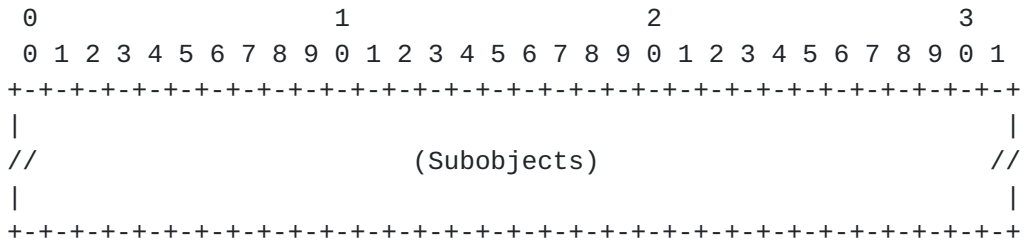

The contents of a PRIMARY_PATH_ROUTE object are a series of variable-length data items called subobjects. The subobjects are identical to those that can constitute an EXPLICIT ROUTE object as defined in [RFC-3209], [RFC-3473] and [RFC-3477].

A Path message may contains multiple PRIMARY_PATH_ROUTE objects, where each object is meaningful. This is useful when a given secondary LSP must be link/node/SRLG disjoint from more than one primary LSP (i.e. is protecting more than one primary LSP).

14.2 Applicability

The PRIMARY_PATH_ROUTE object is to be used only when all GMPLS nodes along the path support the PRIMARY_PATH_ROUTE object. The PRIMARY_PATH_ROUTE object is assigned a class value of the form 0bbbbbbb. GMPLS nodes along the path that do not support this object MUST respond with an "Unknown Object Class" error.

14.3 Subobjects

The contents of a PRIMARY_PATH_ROUTE object is identical to the EXPLICIT ROUTE object of the primary LSP and thus defined as a list of variable-length data items called subobjects. Each 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.

As for the EXPLICIT ROUTE object, 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 Interfaces (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 system SHOULD return a "Bad PRIMARY PATH_ROUTE object" error.

14.4 Procedures

TBD.

15. Security Considerations

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

16. Acknowledgments

17. IANA Considerations

IANA assigns values to RSVP protocol parameters. Within the current document a PROTECTION object 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)

18. 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.

19. References

19.1 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-03.txt, August 2002.

[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-07, October 2002.

[LMP-WDM] A.Fredette and J.Lang (Editors), "Link Management Protocol (LMP) for DWDM Optical Line Systems," Internet Draft, Work in progress, draft-ietf-ccamp-lmp-wdm-01.txt, September 2002.

[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.

[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-01.txt, November 2002.

19.2 Informative References

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

20. Author's Addresses

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

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

Full Copyright Statement