CCAMP Working Group
Internet Draft
Expiration Date: June 2004

CY Lee
A. Farrel
S. De Cnodder
December 2003

Exclude Routes - Extension to RSVP-TE

<draft-ietf-ccamp-rsvp-te-exclude-route-01.txt>

1. Status of this memo

2. Abstract

The current RSVP-TE specification, "RSVP-TE: Extensions to RSVP for LSP Tunnels" (RFC 3209) and GMPLS extensions to RSVP-TE, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions" (RFC 3473) allow abstract nodes and resources to be explicitly included in a path setup, but not to be explicitly excluded.

In some networks where precise explicit paths are not computed at the head end it may be useful to specify and signal abstract nodes and resources that are to be explicitly excluded from routes. These exclusions may apply to the whole path, or to parts of a path between two abstract nodes specified in an explicit path. How Shared Risk Link Groups (SLRGs) can be excluded is also specified in this document.

This document specifies ways to communicate route exclusions during path setup using RSVP-TE.

2.1 Future Work

Future work on this document may include the following.

- Exclusion of unnumbered links.

- Line up with LSP attribute. This could mean that EXRS has to be revised.

- Convergence of SRLG identification with formats defined in other drafts.

- Update MIB section.

2.2 Changes compared to version 00

- This section is added.

- Tolerance field in SRLG subobject is removed.

- References updated.

- Editorial updates.

- XRO processing rules further detailed.

- Recommendation added to limit the size of the exlude route list to a value local to the node originating the exclude route list.

- Section added with minimum compliance statement.

- Acknowledgements updated.

- IPR section.

- Appendix A with applications is added.

3. Conventions used in this document

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

4. Overview

The current RSVP-TE specification [RSVP-TE] and GMPLS extensions [GMPLS-RSVP-TE] allow abstract nodes and resources to be explicitly

included in a path setup, using the Explicit Route Object (ERO).

In some systems it may be useful to specify and signal abstract nodes and resources that are to be explicitly excluded from routes. This may be because loose hops or abstract nodes need to be prevented from selecting a route through a specific resource. This is a special case of distributed path calculation in the network.

Two types of exclusions are required:

i) Exclude any of the abstract nodes in a given set anywhere on the path. This set of abstract nodes is referred to as the Exclude Route list.

ii) Exclude certain abstract nodes or resources between a specific pair of abstract nodes present in an ERO. Such specific exclusions are referred to as Explicit Exclusion Route.

To convey these constructs within the signaling protocol, a new RSVP object and a new ERO subobject are introcuded respectively.

i) A new RSVP-TE object is introduced to convey the Exclude Route list. This object is the Exclude Route Object (XRO).

ii) The second type of exclusion is achieved through a modification to the existing ERO. A new subobject type the Explicit Exclude Route Subobject (EXRS) is introduced to indicate an exclusion between a pair of included abstract nodes.

SRLGs allow the definition of resources or groups of resources that share the same risk of failure. The knowledge of SRLGs may be used to compute diverse paths that can be used for protection. In systems where it is useful to signal exclusions, it may be useful to signal SRLGs to indicate groups of resources that should be excluded on the whole of a path or between two abstract nodes specified in an explicit path.

This document introduces an ERO subobject to indicate an SRLG to be signaled in either of the two exclusion methods described above. This subobject might also be appropriate for use within Explicit Routes or Record Routes, but that discussion is outside the scope of this document.

4.1 Scope of Exclude Routes

This document does not preclude a route exclusion from listing many nodes or network elements to avoid. The intent is, however, to indicate only the minimal number of subobjects to be avoided. For

instance it may be necessary to signal only the SRLGs (or Shared Risk Groups) to avoid.

It is envisaged that most of the conventional inclusion subobjects are specified in the signaled ERO only for the area where they are pertinent. The number of subobjects to be avoided, specified in the signaled XRO may be constant throughout the whole path setup, or the subobjects to be avoided may be removed from the XRO as they become irrelevant in the subsequent hops of the path setup.

For example, consider an LSP that traverses multiple computation domains. A computation domain may be an area in the administrative or IGP sense, or may be an arbitrary division of the network for active management and path computational purposes. Let the primary path be (Ingress, A1, A2, AB1, B1, B2, BC1, C1, C2, Egress) where:

- Xn denotes a node in domain X, and

- XYn denotes a node on the border of domain X and domain Y.

Note that Ingress is a node in domain A, and Egress is a node in domain C. This is shown in Figure 1 where the domains correspond with areas.

```
             area A              area B              area C
     <-------------------> <----------------> <------------------->

 Ingress-----A1----A2----AB1----B1----B2----BC1----C1----C2----Egress
    ^  \                / | \              / | \                /
    |   \              /  |  \            /  |  \              /
    |    A3----------A4--AB2--B3--------B4--BC2--C3----------C4
    |                     ^                  ^
    |                     |                  |
    |                     |              ERO: (C3-strict, C4-strict,
    |                     |                    Egress-strict)
    |                     |              XRO: Not needed
    |                     |
    |          ERO: (B3-strict, B4-strict, BC2-strict, Egress-loose)
    |          XRO: (C1, C2)
    |
   ERO: (A3-strict, A4-strict, AB2-strict, Egress-loose)
   XRO: (B1, B2, BC1, C1, C2, Egress)
```

Consider the establishment of a node-diverse protection path in the example above. The protection path must avoid all nodes on the primary path. The exclusions for area A are handled during Constrained

Shortest Path First (CSPF) computation at Ingress, so the ERO and XRO signaled at Ingress could be (A3-strict, A4-strict, AB2-strict, Egress-loose) and (B1, B2, BC1, C1, C2) respectively. At AB2 the ERO and XRO could be (B3-strict, B4-strict, BC2-strict, Egress-loose) and (C1,C2) respectively. At BC2 the ERO could be (C3-strict, C4-strict, Egress-strict) and an XRO is not needed from BC2 onwards.

In general, consideration should be given (as with explicit route) to the size of signaled data and the impact on the signaling protocol.

4.2 Relationship to MPLS TE MIB

[MPLS-TE-MIB] defines managed objects for managing and modeling MPLS-based traffic engineering. Included in [MPLS-TE-MIB] is a means to configure explicit routes for use on specific LSPs. This configuration allows the exclusion of certain resources.

In systems where the full explicit path is not computed at the ingress (or at a path computation site for use at the ingress) it may be necessary to signal those exclusions. This document offers a means of doing this signaling.

5. Shared Risk Link Groups

The identifier of a SRLG is defined as a 32 bit quantity in [GMPLS-OSPF]. These 32 bits are divided into an 8 bit type field and a 24 bit identifier in [CCAMP-SRLG].

5.1 SRLG ERO Subobject

The format of the ERO and its subobjects are defined in [RSVP-TE]. The new SRLG subobject is defined by this document 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    |       SRLG Id (4 bytes)       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      SRLG Id (continued)      |           Reserved            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L

The L bit is an attribute of the subobject. The L bit is set if the subobject represents a loose hop in the explicit route. If the bit is not set, the subobject represents a strict hop in the explicit route.

For exclusions, the L bit SHOULD be set to zero and ignored.

Type

The type of the subobject [TBD].

Length

The Length contains the total length of the subobject in bytes, including the Type and Length fields. The Length is always 8.

SRLG Id

The 32 bit identifier of the SRLG.

Reserved

Zero on transmission. Ignored on receipt

6. Exclude Route List

The exclude route identifies a list of abstract nodes that MUST NOT be traversed along the path of the LSP being established. It is RECOMMENDED to limit size of the exlude route list to a value local to the node originating the exclude route list.

6.1 Exclude Route Object (XRO)

Abstract nodes to be excluded from the path are specified via the EXCLUDE_ROUTE object (XRO). The Exclude Route Class value is [TBD].

Currently one C_Type is defined, Type 1 Exclude Route. The EXCLUDE_ROUTE object has the following format:

Class = TBD, 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)                         //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Subobjects

The contents of an EXCLUDE_ROUTE object are a series of variable-length data items called subobjects. The subobjects are identical

to those defined in [RSVP-TE] and [GMPLS-RSVP-TE] for use in EROs.

The following subobject types are supported.

Type	Subobject
1	IPv4 prefix
2	IPv6 prefix
32	Autonomous system number
TBD	SRLG

The defined values for Type above are specified in [RSVP-TE] and in this document.

The concept of loose or strict hops has no meaning in route exclusion. The L bit, defined for ERO subobjects in [RSPV-TE], is re-used here to indicate that an abstract node MUST be avoided (value 0) or SHOULD be avoided (value 1).

An Attribute octet is introduced in the subobjects that define IP addresses to indicate the attribute (e.g. interface, node, SRLG) associated with the IP addresses that can be excluded from the path. For instance, the attribute node allows a whole node to be excluded from the path, in contrast to the attribute interface, which allows specific interfaces to be excluded from the path. The attribute SRLG allows all SRLGs associated with an IP address to be excluded from the path.

6.1.1 Subobject 1: IPv4 prefix

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    | IPv4 address (4 bytes)        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv4 address (continued)      | Prefix Length |   Attribute   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L

0 indicates that the attribute specified MUST be excluded
1 indicates that the attribute specified SHOULD be avoided

Attribute

interface

0 indicates that the interface or set of interfaces associated with the IP prefix should be excluded or avoided

node

1 indicates that the node or set of nodes associated with the IP prefix should be excluded or avoided

SRLG

2 indicates that all the SRLGs associated with the IP prefix should be excluded or avoided

Resvd
Zero on transmission. Ignored on receipt.

The rest of the fields are as defined in [RSVP-TE].

6.1.2 Subobject 2: IPv6 Prefix

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    | IPv6 address (16 bytes)       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv6 address (continued)                                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv6 address (continued)                                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv6 address (continued)                                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv6 address (continued)      | Prefix Length |   Attribute   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

L
0 indicates that the abstract node specified MUST be excluded
1 indicates that the abstract node specified SHOULD be avoided

Attribute

interface

0 indicates that the interface or set of interfaces associ- ated with the IP prefix should be excluded or avoided

node

1 indicates that the node or set of nodes associated with the IP prefix should be excluded or avoided

SRLG

2 indicates that all the SRLG associated with the IP prefix should be excluded or avoided

Resvd
Zero on transmission. Ignored on receipt.

The rest of the fields are as defined in [RSVP-TE].

6.1.3 Subobject 32: Autonomous System Number

The L bit of an Autonomous System Number subobject has meaning in an Exclude Route (contrary to its usage in an Explict Route defined in [RSVP-TE]. The meaning is as for other subobjects described above. That is:

0 indicates that the abstract node specified MUST be excluded

1 indicates that the abstract node specified SHOULD be avoided

The rest of the fields are as defined in [RSVP-TE]. There is no Attribute octet defined.

6.1.4 Subobject TBD: SRLG

The Attribute octet is not present. The rest of the fields are as defined in the "SRLG ERO Subobject" section of this document.

6.2. Semantics and Processing Rules for the Exclude Route Object (XRO)

The exclude route list is encoded as a series of subobjects contained in an EXCLUDE_ROUTE object. Each subobject identifies an abstract node in the exclude route list.

Each abstract node may be a precisely specified IP address belonging to a node, or an IP address with prefix identifying interfaces of a group of nodes, or an Autonomous System.

The Explicit Route and routing processing is unchanged from the description in [RSVP-TE] with the following additions:

a. When a Path message is received at a node, the node must check that it is not a member of any of the abstract nodes in the XRO if it is present in the Path message. If the node is a member of any of the abstract nodes in the XRO with the L-flag set to "exclude", it should return a PathErr with the error code "Routing Problem" and error value of "Local node in Exclude Route". If there are SRLGs in the XRO, the node should check that the resources the node uses are not part of any SRLG with the L-flag set to

"exclude" that is specified in the XRO. If it is, it should return a PathErr with error code "Routing Problem" and error value of "Local node in Exclude Route".

b. Each subobject must be consistent. If a subobject is not consistent then the node should return a PathErr with error code "Routing Problem" and error value "Inconsistent Subobject". An example of an inconsistent subobject is an IPv4 Prefix subobject containing the IP address of a node and the attribute field is set to "interface" or "SRLG".

c. The subobjects in the ERO and XRO SHOULD not contradict each other. If they do contradict, the subobjects with the L flag not set, strict or MUST be excluded, respectively, in the ERO or XRO MUST take precedence. If there is still a conflict, a PathErr with error code "Routing Problem" and error value of "Route blocked by Exclude Route" should be returned.

d. When choosing a next hop or expanding an explicit route to include additional subobjects, a node:

   i) must not introduce an explicit node or an abstract node that equals or is a member of any abstract node that is specified in the Exclude Route Object with the L-flag set to "exclude". The number of introduced exlicit nodes or abstract nodes with the L flag set to "avoid" should be minimised.

   ii) must not introduce links, nodes or resources identified by the SRLG Id specified in the SRLG subobjects(s). The number of introduced SLRGs with the L flag set to "avoid" should be minimised.

   If these rules preclude further forwarding of the Path message, the node should return a PathErr with the error code "Routing Problem" and error value of "Route blocked by Exclude Route".

   Note that the subobjects in the XRO is an unordered list of subobjects.

The XRO Class-Num is of the form 11bbbbbb so that nodes which do not support the XRO will forward it uninspected and will not apply the extensions to ERO processing described above. This makes the XRO a 'best effort' process.

This 'best-effort' approach is chosen to allow route exclusion to

traverse parts of the network that are not capable of parsing or handling the new function. Note that Record Route may be used to allow computing nodes to observe violations of route exclusion and attempt to re-route the LSP accordingly.

If a node supports the XRO, but not a particular subobject or part of that subobject, then that particular subobject is ignored. Examples of a part of a subobject that can be supported are: (1) only prefix 32 of the IPv4 prefix subobject could be supported, or (2) a particular subobject is supported but not the particular attribute field.

When a node forwards a Path message, it can do the following three operations related to XRO besides of the processing rules mentioned above:

1. If no XRO was present, an XRO may be included.

2. If an XRO was present, it may remove the XRO if it is sure that the next nodes do not need this information anymore. An example is where a node can expand the ERO to a full strict path towards the destination. See Figure 1 where BC2 is removing the XRO from the Path message.

3. If an XRO was present, the content of the XRO can be modified. Subobjects can be added or removed. See Figure 1 for an example where AB2 is stripping off some subobjects.

7. Explicit Exclude Route

The Explicit Exclude Route defines abstract nodes or resources (such as links, unnumbered interfaces or labels) that must not be used on the path between two inclusive abstract nodes or resources in the explicit route.

7.1. Explicit Exclusion Route Subobject (EXRS)

A new ERO subobject type is defined. The Explicit Exclude Route Subobject (EXRS) has type [TBD]. The EXRS may not be present in an RRO or XRO.

The format of the EXRS is as follows.

```
  0                   1
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+--------------//---------------+
 |L|    Type     |     Length    |     EXRS subobjects           |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+--------------//---------------+
```

L

ignored and must be zero
[Note: The L bit in an ERES subobject is as defined for the XRO subobjects]

Type

The type of the subobject, i.e. EXRS [TBD]

EXRS subobjects

An EXRS subobject indicates the abstract node or resource to be excluded. The format of this field is exactly the format of an XRO subobject and may include an SRLG subobject. Both subobjects are as described earlier in this document.

Thus, an EXRO subobject for an IP hop might look 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|L|    Type     |     Length    |L|    Type     |     Length    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| IPv4 address (4 bytes)                                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Prefix Length |   Attribute   |          Reserved             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Note: The Most Significant Bit in the Type field could be used to indicate exclusion of IPv4/IPv6, AS and SRLG subobjects, eliminating the need to prepend the subobject with an additional TLV header. This would reduce the number bytes require for each subobject by 2 bytes. However, this approach would reduce the ERO Type field space by half. This issue need WG discussion and feedback.

7.2. Semantics and Processing Rules for the EXRS

Each EXRS may carry multiple exclusions. The exclusion is encoded exactly as for XRO subobjects and prefixed by an additional Type and Length.

The scope of the exclusion is the step between the previous ERO subobject that identifies an abstract node, and the subsequent ERO subobject that identifies an abstract node. Multiple exclusions may be present between any pair of abstract nodes.

Exclusions may indicate explicit nodes, abstract nodes or Autonomous Systems that must not be traversed on the path to the next abstract node indicated in the ERO.

Exclusions may also indicate resources (such as unnumbered interfaces, link ids, labels) that must not be used on the path to the next abstract node indicated in the ERO.

SRLGs may also be indicated for exclusion from the path to the next abstract node in the ERO by the inclusion of an EXRO Subobject containing an SRLG subobject. If the Tolerance value in the SRLG subobject is zero, the resources (nodes, links, etc.) identified by the SRLG must not be used on the path to the next abstract node indicated in the ERO. If the Tolerance value is non- zero, the resources identified by the SRLG should be avoided, but may be used in preference to resources associated with another SRLG indicated for exclusion if that SRLG has a (numerically) lower Tolerance value.

The subobjects in the ERO and EXRS SHOULD not contradict each other. If they do contradict, the subobjects with the L bit not set, strict or MUST be excluded, respectively, in the ERO or XRO MUST take precedence. If there is still a conflict, the subobjects in the ERO MUST take precedence.

If a node is called upon to process an EXRS and does not support handling of exclusions it will return a PathErr with a "Bad EXPLICIT_ROUTE object" error.

If the presence of EXRO Subobjects precludes further forwarding of the Path message, the node should return a PathErr with the error code "Routing Problem" and error value of "Route blocked by Exclude Route".

8. Minimum compliance

An implementation must be at least compliant with the following:

A. The XRO MUST be supported with the following restrictions:

A.1. The IPv4 Prefix subobject MUST be supported with a prefix length of 32, and an attribute value of "interface" and "node". Other prefix values and attribute values MAY be supported.

A.2. The IPv6 Prefix subobject MUST be supported with a prefix length of 128, and an attriubute value of "interface" and "node". Other prefix values and attribute values MAY be supported.

B. The EXRS SHOULD be supported. If supported, the same restrictions as for the XRO apply.

C. If XRO or EXRS are supported, the implementation MUST be compliant

with the processing rules of the supported, not supported, or partially supported subobjects as specified within this document.

9. Security

The new exclude route object poses no security exposures over and above [RSVP-TE] and [GMPLS-RSVP-TE]. Note that any security concerns that exist with Explicit Routes should be considered with regard to route exclusions.

10. IANA Considerations

10.1. New Class Numbers

One new class number is required.

EXCLUDE_ROUTE
Class-Num = 011bbbbb
CType: 1

10.2. New Subobject Types

A new subobject type for the Exclude Route Object and Explicit Exclude Route Subobject is required.

SRLG subobject

A new subobject type for the ERO is required.

Explicit Exclude Route subobject

10.3. New Error Codes

New error values are needed for the error code 'Routing Problem'.

Unsupported Exclude Route Subobject Type	[TBD]
Inconsistent Subobject	[TBD]
Local Node in Exclude Route	[TBD]
Route Blocked by Exclude Route	[TBD]

11. Acknowledgments

This document reuses text from [RSVP-TE] for the description of EXCLUDE_ROUTE.

The authors would like to express their thanks to Lou Berger, Steffen Brockmann, Igor Bryskin, Dimitri Papadimitriou, Cristel Pelsser, and Richard Rabbat for their considered opinions on this draft. Also

thanks to Yakov Rekhter for reminding us about SRLGs!

12. Intellectual Property Considerations

This section is taken from Section 10.4 of [RFC-2026].

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.

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

13. References

13.1 Normative References

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

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

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

[GMPLS-OSPF] K. Kompela, Y. Rekhter, (Editors) "OSPF Extensions in Support of Generalized MPLS", Internet Draft,

draft-ietf-ccamp-ospf-gmpls-extensions-12.txt, October 2003 (work in progress).

[MPLS-TE-MIB] C. Srinivasan, et al., "Multiprotocol Label Switching (MPLS) Traffic Engineering Management Information Base", Internet Draft, draft-ietf-mpls-te-mib-09.txt, November 2002 (work in progress).

13.2 Informational References

[MPLS-BUNDLE] Kompella, K., Rekhter, Y., and Berger, L., "Link Bundling in MPLS Traffic Engineering", Internet Draft, draft-ietf-mpls-bundle-04.txt, July 2002, (work in progress).

[MPLS-UNNUM] Kompella, K., Rekhter, Y., "Signalling Unnumbered Links in RSVP-TE", RFC 3477, January 2003.

[CCAMP-SRLG] D. Papadimitriou, et al., "Shared Risk Link Groups Encoding and Processing", Internet Draft, draft-papadimitriou-ccamp-srlg-processing-01.txt, November 2002 (work in progress).

[INTERAS] De Cnodder, S., Pelsser, C., "Protection for inter-AS MPLS tunnels", Internet Draft, draft-decnodder-mpls-interas-protection-00.txt, February 2003, (work in progress).

[OVERLAY] Swallow, G., Drake, J., Ishimatsu, H., Rekhter, Y., GMPLS RSVP Support for the Overlay Model", Internet Draft, draft-ccamp-gmpls-overlay-02.txt, October 2003, (work in progress).

[OSPF-TE] Katz, D., Yeung, D., and Kompella, K., "Traffic Engineering Extensions to OSPF version 2", RFC 3630, September 2003.

[ISIS-TE] Smit, H., Li, T., "IS-IS extensions for Traffic Engineering", Internet Draft, draft-ietf-isis-traffic-05.txt, August 2003, (work in progress).

[CRANKBACK] Farrel, A., (Editor), "Crankback Routing Extensions for MPLS Signaling", Internet Draft, draft-iwata-mpls-crankback-05.txt, March 2003, (work in progress).

14. Authors' Information

Cheng-Yin Lee
Alcatel
600 March Road.
Ottawa, Ontario
Canada K2K 2E6
email: Cheng-Yin.Lee@alcatel.com

Adrian Farrel
Movaz Networks, Inc.
7926 Jones Branch Drive, Suite 615
McLean VA, 22102 USA
Phone: +1-703-847-1867
Email: afarrel@movaz.com

Stefaan De Cnodder
Alcatel
Francis Wellesplein 1
B-2018 Antwerp, Belgium
email: stefaan.de_cnodder@alcatel.be

15. Appendix A: applications

This section describes some applications that can make use of the XRO. The intention is to show that the XRO is not an application specific object, but that it can be used for multiple purposes. In a few examples, other solutions might be possible for that particular case but the intention is to show that also a single object can be used for all the examples, hence making the XRO a rather generic object without having to define a solution and new objects for each new application.

15.1 Inter-area LSP protection

One method to establish an inter-area LSP is where the ingress router selects an ABR, and then the ingress router computes a path towards this selected ABR such that the configured constraints of the LSP are fulfilled. In the example of figure A.1, an LSP has to be established from node A in area 1 to node C in area 2. If no loose hops are configured, then the computed ERO at A could looks as follows: (A1-strict, A2-strict, ABR1-strict, C-loose). When the Path message arrives at ABR1, then the ERO is (ABR1-strict, C-loose) and it can be expanded by ABR1 to (B1-strict, ABR3-strict, C-loose). Similar, at ABR3 the received ERO is (ABR3-strict, C-loose) and it can be expanded to (C1-strict, C2-strict, C-strict). If also a backup LSP has to be established, then A takes another ABR (ABR2 in this case) and computes a path towards this ABR that fulfills the constraints of the LSP and such that is disjoint from the path of the primary LSP. The ERO generated by A looks as follows for this example: (A3-strict,

A4-strict, ABR2-strict, C-loose).

In order to let ABR2 expand the ERO, it also needs to know the path of the primary LSP to expand the ERO such that it is disjoint from the path of the primary LSP. Therefore, A also includes an XRO that at least contains (ABR1, B1, ABR3, C1, C2). Based on these constraints, ABR2 can expand the ERO such that it is disjoint from the primary LSP. In this example, the ERO computed by ABR2 would be (B2-strict, ABR4-strict, C-loose), and the XRO generated by B contains at least (ABR3, C1, C2). The latter information is needed to let ABR4 to expand the ERO such that the path is disjoint from the primary LSP in area 2.

```
     Area 1           Area 0          Area 2
<---------------><--------------><--------------->

+---A1---A2----ABR1-----B1-----ABR3----C1---C2---+
|        |             |              |         |
|        |             |              |         |
A        |             |              |         C
|        |             |              |         |
|        |             |              |         |
+---A3---A4----ABR2-----B2-----ABR4----C3---C4---+
```

Figure A.1: Inter-area LSPs

In this example, a node performing the path computation, first selects an ABR and then it computes a strict path towards this ABR. For the backup LSP, all nodes of the primary LSP in the next areas has to be put in the XRO (with the exception of the destination node if node protection and no link protection is required). When an ABR computes the next path segment, i.e. the path over the next area, it may remove the nodes from the XRO that are located in that area with the exception of the ABR where the primary LSP is exiting the area. The latter information is still required because when the selected ABR (ABR4 in this example) further expands the ERO, it has to exclude the ABR on which the primary is entering that area (ABR3 in this example). This means that when ABR2 generates an XRO, it may remove the nodes in area 0 from the XRO but not ABR3. Note that not doing this would not harm in this example because there is no path from ABR4 to C via ABR3 in area2. If there would be a links between ABR4-ABR3 and ABR3-C, then it is required to have ABR3 in the XRO generated by ABR2.

Discussion on the length of the XRO: when link or node protection is requested, the length of the XRO is bounded by the length of the RRO of the primary LSP. It can be made shorter by removing nodes by the ingress node and the ABRs. In the example above, the RRO of the

primary LSP contains 8 subobjects, while the maximum XRO length can be bounded by 6 subobjects (nodes A1 adn A2 do not have to be in the XRO. For SRLG protection, the XRO has to list all SRLGs that are crossed by the primary LSP.

15.2 Inter-AS LSP protection

When an inter-AS LSP is established, which has to be protected by a backup LSP to provide link or node protection, the same method as for the inter-area LSP case can be used. The difference is when the backup LSP is not following the same AS-path as the primary LSP because then the XRO should always contain the full path of the primary LSP. In case the backup LSP is following the same AS-path (but with different ASBRs - at least in case of node protection), it is much similar as the inter-area case: ASBRs expanding the ERO over the next AS may remove the XRO subobjects located in that AS. Note that this can only be done by ingress ASBRs (the ASBR where the LSP is entering the AS).

Discussion on the length of the XRO: the XRO is bounded by the length of the RRO of the primary LSP.

Suppose that SRLG protection is required, and the ASs crossed by the main LSP use a consistent way of allocating SRLG-ids to the links (i.e. the ASs use a single SRLG space). In this case, the SRLG-ids of each link used by the main LSP can be recorded by means of the RRO, which are then used by the XRO. If the SRLG-ids are only meaningfull local to the AS, putting SRLG-ids in the XRO crossing many ASs makes no sense. More details on the method of providing SRLG protection for inter-AS LSPs can be found in [INTERAS]. Basically, the link IP address of the inter-AS link used by the primary LSP is put into the XRO of the Path message of the detour LSP or bypass tunnel. The ASBR where the detour LSP or bypass tunnel is entering the AS can translate this into the list of SRLG-ids known to the local AS.

Discussion on the length of the XRO: the XRO only contains 1 subobject, which contains the IP address of the inter-AS link traversed by the primary LSP (in the assumption that the primary LSP and detour LSP or bypass tunnel are leaving the AS in the same area, and they are also entering the next AS in the same area).

15.3 Protection in the GMPLS overlay model

When an edge-node wants to establish an LSP towards another edge-node over an optical core network as described in [OVERLAY] (see figure A.2), the XRO can be used for multiple purposes.

```
    Overlay                                                    Overlay
    Network        +----------------------------------+        Network
  +----------+     |                                  |     +----------+
  |   +----+ |     |  +-----+    +-----+    +-----+   |     | +----+   |
  |   |    | |     |  |     |    |     |    |     |   |     | |    |   |
  | --+ EN1+-+-----+--+ CN1 +----+ CN2 +----+ CN3 +---+-----+-+ EN3+-- |
  |   |    | |  +--+--+     |    |     |    |     +---+--+  | |    |   |
  |   +----+ |  |  |  +--+--+    +--+--+    +--+--+   |  |  | +----+   |
  |          |  |  |     |          |          |      |  |  |          |
  +----------+  |  |     |          |          |      |  |  +----------+
                |  |     |          |          |      |  |
  +----------+  |  |     |          |          |      |  |  +----------+
  |          |  |  |  +--+--+       |       +--+--+   |  |  |          |
  |   +----+ |  |  |  |     |       +-------+     |   |  |  | +----+   |
  |   |    +-+--+  |  | CN4 +---------------+ CN5 |   |  +--+-+    |   |
  | --+ EN2+-+-----+--+     |               |     +---+-----+-+ EN4+-- |
  |   |    | |     |  +-----+               +-----+   |     | |    |   |
  |   +----+ |     |                                  |     | +----+   |
  |          |     +----------------------------------+     |          |
  +----------+                  Core Network                +----------+
    Overlay                                                    Overlay
    Network                                                    Network

                                             Legend:   EN  -  Edge Node
                                                       CN  -  Core  Node
```

Figure A.2

A first application is where an edge-node wants to establish multiple LSPs towards the same destinatin edge-node, and these LSPs need to have as few or no SRLGs in common. In this case EN1 could establish an LSP towards EN3 and then it can establish a second LSP listing all links used by the first LSP with the indicition to avoid the SRLGs of these links. This information can be used by CN1 to compute a path for the second LSP. If the core network consists of multiple areas, then the SRLG-ids have to be listed in the XRO. The same example applies to nodes and links.

Another application is where the edge-node wants to set up a backup LSP that is also protecting the links between the edge-nodes and core-nodes. For instance, when EN2 establishes an LSP to EN4, it sends a Path message to CN4, which computes a path towards EN4 over for instance CN5. When EN2 gets back the RRO of that LSP, it can signal a new LSP to CN1 with EN4 as destination and the XRO computed based on the RRO of the first LSP. Based on this information, CN1 can compute a path that has the requested diversaty properties (e.g, a path going over CN2, CN3 and then to EN4).

It is clear that in these examples, the core-node may not edit the

RRO in a Resv message such that it includes only the subobjects from the egress core-node through the egress edge-node.

15.4 LSP protection inside a single area

The XRO can also be used inside a single area. Take for instance a network where the TE extensions of the IGPs as described in [OSPF-TE] and [ISIS-TE] are not used, and hence each node has to select a next-hop and possibly crankback [CRANKBACK] has to be used when there is no viable next-hop. In this case, when signaling a backup LSP, the XRO can be put in the Path message to exclude the links, nodes or SRLGs of the primary LSP. An alternative to provide this functionality would be to indicate in the Path message of the backup LSP, the primary LSP together witn an indication which type of protection is required. This latter solution would work for link and node protection, but not for SRLG protection.

Discussion on the length of the XRO: when link or node protection is requested, the XRO is of the same length as the RRO of the primary LSP. For SRLG protection, the XRO has to list all SRLGs that are crossed by the primary LSP. Note that for SRLG protection, the link IP address to reference the SRLGs of that link cannot be used since the TE extensions of the IGPs are not used in this example, hence, a node cannot translate any link IP address located in that area to its SRLGs.

16. Full Copyright Statement