

September 2003

Protection for inter-AS MPLS tunnels
[draft-decnodder-mpls-interas-protection-01.txt](#)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

Abstract

This document describes a solution for link protection, node protection, Shared Risk Link Group (SRLG) protection and fast recovery of inter-AS LSPs. These problems are highlighted in [\[ASREQ\]](#). The proposed solution is based on RSVP-TE [\[RFC3209\]](#) as recommended by [\[ASREQ\]](#).

[1. Introduction](#)

This document describes a solution for the following requirements from [\[ASREQ\]](#):

- 1) link protection
- 2) node protection

- 3) SRLG protection
- 4) fast recovery
- 5) based on RSVP-TE [[RFC3209](#)]

MPLS Fast-Reroute techniques based on [[FRR](#)] together with the RSVP objects Exclude Route Object (XRO) and Explicit Exclude Route Subobject (EXRS), as defined in [[XRO](#)], will be used to fulfill the above requirements. Only the protection of links between 2 ASs, the protection of their SRLGs and the protection of nodes at the border of an AS are in the scope of this document.

[Section 3](#) proposes to tunnel inter-AS LSPs through intra-AS LSPs inside an AS, as described in [[HIER](#)]. This tunneling favors the confidentiality requirement concerning intra-AS topologies [[ASREQ](#)] as well as the establishment of inter-AS LSPs. The establishment of inter-AS LSPs will not be studied further in this draft. In this document it is assumed that ASes define their SRLGs independent from the SRLGs in other ASes.

[Section 4](#) shows that an end-to-end backup LSP can only provide link and node protection. For SRLG protection and fast recovery, the methods in [[FRR](#)] have to be used. These methods are described in [section 6](#) for detour LSPs and the use of bypass tunnels to protect inter-AS LSPs is introduced in [section 7](#). Nodes other than those mentioned in this document, must use the methods in [[FRR](#)] to establish detour LSPs or bypass tunnels. This means that these nodes establish detour LSPs that merge with the main LSP in the same AS where they are originated, or these nodes establish bypass tunnels that terminate in the same AS as where they originate.

2. Summary for Sub-IP Area

<TBD>

3. Inter-AS LSP tunneled through an intra-AS LSP

To improve scalability and confidentiality (which is outside the scope of this document), an inter-AS LSP can be tunneled through an intra-AS LSP [[HIER](#)]. For instance, in Figure 2, the link between R23 and R24 could be an LSP passing multiple core routers. And, the inter-AS LSP is tunneled through this LSP.

The procedures described in the following sections apply for inter-AS link, node and SRLG protection of inter-AS LSPs whether they are tunneled or not.

4. Problems to protect SRLGs with disjoint end-to-end LSPs

The motivation to support fast-reroute techniques as described in [FRR] is twofold: first of all, it supports fast recovery and secondly, it can also provide SRLG protection, which is not the case for a disjoint end-to-end LSP. The problems to support SRLG protection, with the latter method, are described in this section.

There are different ways to provide end-to-end protection of inter-AS LSPs. A first possibility is to establish a secondary path that crosses different ASs than the main LSP. An alternative is to establish an LSP that follows the same AS path to the destination as the main LSP, i.e. it crosses the same ASs in the same order, but is link or node disjoint from the main LSP. However, these two solutions do not permit to establish an LSP that is disjoint from the SRLGs of the main LSP. That is, it is not possible to protect the main inter-AS LSP against SRLGs failures with a single end-to-end link or node disjoint LSP. This is due to the fact that ASs may possess links belonging to the same SRLG even if these ASs do not have the same convention to designate this SRLG.




Figure 1: end-to-end SRLG protection

For example, on figure 1, we have a main LSP going from R11 in AS1 to R33 in AS3 through R13, R22, R24 and R32. It is not possible to protect this LSP against SRLG failures with a backup LSP crossing R12, R21, R23 and R31. This is because AS3 could have links which have SRLGs in common with links in AS1 and AS1 nor AS3 will be aware of it. For example, link R11-R13 and link R31-R33 may belong to the same SRLG. This example relies on the fact that different ASs may use the same resources to join different nodes in their respective domain. A similar situation occurs when the main and the backup LSP do not share the same AS path but instead partially cross different ASs.

This document only focuses on local protection as defined in [FRR],

because it is not possible to provide full protection of an inter-AS LSP with a single end-to-end LSP. The solution proposed in this document enables the provision of link, node and SRLG protection of inter-AS LSPs.

5. Network model and terminology

To illustrate the procedures described in the next sections, the following network model is used:




Figure 2: a reference network model

The main LSP is established from a certain node (not shown on the figure) and goes over routers R13, R14, R23, and R24 towards the destination (also not shown on the figure). AS1 is referred to as the upstream AS of AS2, and AS2 is referred to as the downstream AS of AS1.

An "egress AS-BR" or a "primary egress AS-BR" is an Autonomous System Border Router (AS-BR) at which the main LSP leaves an AS. In the network example, in figure 2, this is router R14, inside AS1.

An "ingress AS-BR" or a "primary ingress AS-BR" is an AS-BR at which the main LSP enters an AS. In the network example, this is router R23, inside AS2.

A "secondary egress AS-BR" is an AS-BR at which the bypass tunnel or the detour LSP leaves an AS. In the network example, this could be router R12, in AS1.

A "secondary ingress AS-BR" is an AS-BR at which the bypass tunnel or the detour LSP enters an AS. In the network example, this could be router R21, in AS2.

"Inter-AS link protection" is the protection of an LSP against a failure of the link connecting two ASs on the path of the LSP. In the network example, the inter-AS link R14-R23 is to be protected.

"Inter-AS node protection" is the protection of an LSP against an AS-BR failure. This can be the egress AS-BR, R14, or the ingress AS-BR, R23, for the considered example.

"Inter-AS SRLG protection" is the protection of an LSP against a simultaneous failure of all links that belong to certain SRLGs which also contain the inter-AS link (R14-R23 in figure 2).

Other terminology and abbreviations are taken from [[FRR](#)].

6. Protection with detour LSPs

6.1 Link protection with detour LSPs

6.1.1 Procedures for the egress AS-BR

The primary egress AS-BR has to establish a detour LSP to protect the interdomain link. The destination of the detour LSP will be the same as the destination of the main LSP. The detour LSP may merge with the main LSP at any downstream node or with other detour LSPs of the same main LSP, established by nodes downstream of the link to be protected. The egress AS-BR has to determine a secondary egress AS-BR and then it can perform a path calculation towards this AS-BR.

The primary egress AS-BR can select any other AS-BR as secondary egress AS-BR but it is recommended to select an AS-BR that is connected to the downstream AS of the main LSP (i.e. the AS where the primary ingress AS-BR is located). In case this condition is not met, it could be for instance possible that the downstream AS of the detour LSP chooses a path that goes through the AS where the detour LSP was originated which can cause loops. In addition, it is recommended that the detour LSP merges in the AS where this downstream ingress AS-BR is located (the merging node could be the ingress AS-BR itself) if the destination of the main LSP is not in the downstream AS (AS2 in Figure 2). This suggestion improves the scalability of the solution since merging of LSPs diminishes the number of states to be maintained, the bandwidth to be reserved, and so on. Therefore, the egress AS-BR should put a portion of the RRO of the main LSP inside the ERO of the detour LSP. The last hop in the downstream AS (the egress AS-BR in the downstream AS) of the main LSP and all hops after that router should be at least in the ERO of the detour LSP. In addition to this portion of the RRO of the main LSP, the ERO may be further prepended by the egress AS-BR with a strict or a loose path towards the selected secondary egress AS-BR. That is,

the ERO of the detour LSP at least contains: (1) A strict or loose path toward the secondary egress ASBR (2) The path of the main LSP starting at the last hop inside the downstream AS and ending at the destination of the main LSP. In the example network, we have a main LSP with RRO containing routers R11, R12, R21, R22, etc. The ERO of the detour LSP protecting link R12-R21 is composed of router R14, R22 (with loose flag set) and the following routers.

There are two possible methods to determine the secondary egress AS-BR at the primary egress AS-BR. (1) The egress AS-BR can be manually configured with other AS-BRs that peer to the same AS or (2) it can lookup in its BGP table to find an other entry such that the AS-path has the same AS next hop as the currently selected entry. Option (1) is feasible because the number of links between 2 ASs is usually limited to only a small number of links.

It could be possible that the primary egress AS-BR is the same router as the secondary egress AS-BR and that the primary ingress AS-BR is the same router as the secondary ingress AS-BR. In this particular case the inter-domain link on the primary path must not be the same link used by the detour LSP and no path calculation should be done to calculate a (partial) path for the detour LSP.

The use of the LSP-Merge subobject, defined in [Appendix A](#), is optional to provide link protection.

6.1.2 Procedures for the ingress AS-BR

No extra procedures are required.

The detour LSP may merge with the main LSP at this node.

6.1.3 Procedures for the secondary egress AS-BR

The secondary egress AS-BR completes the path in the ERO by selecting a secondary ingress AS-BR in the downstream AS. If there is no ERO present, then the tunnel end point address in the Session object has to be used to route the Path message.

6.1.4 Procedures for the secondary ingress AS-BR

The secondary ingress AS-BR completes the ERO with a path towards the next subobject in the ERO. The LSP should merge with the main LSP at the node that processes the LSP-Merge subobject (if that subobject is used), if it was not yet merged at this point. If no ERO is present inside the Path message of the detour LSP, the path is computed based on the tunnel end point address.

6.2 Node protection with detour LSPs

The procedures and recommendations are the same for the protection of an ingress AS-BR failure as for link protection, with the exception that the egress AS-BR has to include an XRO object or an EXRS subobject [[XRO](#)] with the ingress AS-BR to exclude.

For the protection of the egress AS-BR, the same holds except that the procedures apply to the router on the path of the main LSP preceding the egress AS-BR. The method to determine a secondary egress AS-BR is the same as for the egress AS-BR for link protection: either manual configuration or by using BGP routing information, if it is available. Note that the first solution requires more configuration as for link protection in case this router peers with more than one AS-BR.

6.3 SRLG protection with detour LSPs

Similar procedures as for link protection apply for SRLG protection. In addition, the secondary egress AS-BR must be an AS-BR that peers with the downstream AS of the primary LSP. And, the detour LSP must merge at that AS. The former condition is necessary because only the two peering ASs know the SRLGs of the inter-domain link and the latter condition implies that the LSP-Merge subobject must be used. This subobject is inserted inside the ERO to indicate the node where merging needs to be done (see [appendix A](#)). The next subsections describe in more details the procedures to be performed at the nodes involved in the establishment of such detour LSP.

6.3.1 Procedures for the egress AS-BR

The egress AS-BR has to include an XRO object or an EXRS subobject to exclude the SRLGs of the inter-domain link. The XRO or the EXRS must include a list of SRLGs (defined for the AS containing the PLR) corresponding to the inter-AS link as well as a reference to this link. If the egress AS-BR can calculate a strict path to reach the secondary egress AS-BR, then the list of SRLGs may be removed. Only the reference to the link for which the detour LSP has to be SRLG disjoint is then required (see [section 6.3.2](#)). The secondary ingress AS-BR has to use the information in the XRO or EXRS to further calculate a path for the detour LSP.

To ensure merging inside the downstream AS, the LSP-Merge subobject (see [Appendix A](#)) has to be included in the ERO by the egress AS-BR. The LSR where the detour LSP is merged with the main LSP has to ensure that it can perform a switch-over from the incoming detour LSP containing the LSP-Merge subobject to its originating detour LSP in case the next link has an SRLG in common with the inter-domain link.

This is because in this case, both links can fail at the same time such that both detour LSPs will be activated at the same time.

6.3.2 Procedures for the secondary egress AS-BR

The secondary egress AS-BR selects a next hop and the XRO or EXRS contains a reference to the link for which the detour LSP has to be SRLG disjoint. No list of SRLGs should be included because the SRLG IDs are local to an AS, which means that if a list of SRLG IDs would be sent to the next hop, then this node would not understand the IDs. Therefore only the reference to the inter-AS link is useful. This link is referenced by means of its IP address, see [[XRO](#)]. The secondary egress AS-BR thus removes the list of SRLGs related to the inter-AS link, if such a list of SRLGs was present.

6.3.3 Procedures for the ingress AS-BR

No extra procedures required.

6.3.4 Procedures for the secondary ingress AS-BR

If the secondary ingress AS-BR cannot compute a full path towards the node immediately preceding the LSP-merge subobject, then the secondary ingress AS-BR adds the list of SRLGs of the inter-AS link present inside the received XRO or EXRS inside these objects. These SRLGs are known by the nodes inside this AS. This is required because the LSP can cross nodes inside the AS which do not know the SRLGs of the inter-AS link, but only the SRLGs of intra-area links. An alternative would be to distribute inter-AS links and their SRLGs inside the IGP.

6.3.5 Path calculation

To allow the egress AS-BR and the secondary ingress AS-BR to calculate a path, the SRLGs of the inter-AS links towards the same downstream AS (upstream AS, respectively) as the main LSP have to be known. This could be achieved through manual configuration of the SRLGs of other inter-AS links to the same downstream/upstream AS at each AS-BR. For instance, in Figure 2, at R14 and R23, the SRLGs of R12-R21 can be configured such that they are known for the path calculation, and at R12 and R21, the SRLGs of R14-R23 can be configured. An other option is to flood this information via BGP extensions to be defined or to distribute these links and their SRLGs inside the IGP. It is not assumed that nodes other than AS-BRs having a link to the same downstream/upstream AS know the SRLGs of these inter-AS links. If this would be the case, then the procedures above can be simplified, e.g., the egress AS-BR in [Section 6.3.1](#) does not

have to include a list of SRLGs anymore when only a partial path can be computed.

Also the secondary egress AS-BR has to know the SRLGs of the inter-AS link used by the primary LSP. This is to allow the egress AS-BR to select a link in case there are multiple links towards the downstream AS, and to check if the link is indeed SRLG disjoint from the inter-AS link used by the main LSP.

6.3.6 SRLG and node protection

Protection of the egress AS-BR and SRLG protection of the link preceding the egress AS-BR is best solved by using two detour LSPs at the node on the path of the main LSP preceding the egress AS-BR: a detour to protect against the SRLGs of the intra-domain link and a second detour LSP that is established using the procedures for node protection as described in the previous section. The detour protecting against the SRLGs has to merge in the same AS, i.e. it has to merge with the main LSP at the egress AS-BR. This is because other ASs do not know this intra-domain link, nor its SRLGs. To ensure that merging occurs at the egress AS-BR, the RRO of the main LSP should be fully included in the ERO of the detour LSP together with the LSP-Merge subobject. This path should be preceded by a path, which is SRLG disjoint with the next link of the main LSP computed towards the egress AS-BR. This could only be a partial path towards the egress AS-BR in which case an XRO object or an EXRS subobject, containing the SRLGs to avoid, has to be added. It has to be ensured that these 2 detour LSPs do not merge, which means that at least one of the detour LSP should be a sender-template specific detour LSP.

The egress AS-BR must ensure that it can do a switch-over from the incoming detour LSP protecting against a failure of the preceding link to its originating detour LSP. This is because the preceding link and the inter-domain link can belong to the same SRLG, hence they can fail at the same time. For this reason, the LSP-Merge subobject must be used in this case.

If protection of the ingress AS-BR is requested, in addition to SRLG protection, the egress AS-BR also has to put the ingress AS-BR in the XRO or EXRS like it was done for node protection.

The use of 2 detour LSPs (one for SRLG protection and the other for node protection) is also recommended when the ingress AS-BR is to be protected. If only 1 detour LSP is used, and the LSP only crosses 1 hop in the downstream AS (i.e. ingress AS-BR and egress AS-BR in the downstream AS are the same router), the detour LSP setup would fail. This is because the AS further downstream of the immediate downstream AS is not aware anymore of the SRLGs of the link to be protected.

In case of node and SRLG protection or in case of SRLG protection only, it is recommended to use sender-template specific detour LSPs to avoid that detour LSPs merge with each other.

7. Protection with bypass tunnels

The problem of protection by means of bypass tunnels can be split into two parts:

- a) The bypass tunnel has to be signaled over a path that is disjoint with the network resources that it protects.
- b) After the bypass tunnels are established, an appropriate bypass tunnel has to be selected for each particular main LSP such that the protection requirements for that LSP are met.

The first part is very similar to the establishment of detour LSPs: an XRO object or an EXRS subobject can be used to signal the bypass tunnel such that it is disjoint from the network resources used by the main LSP. The same recommendations as for detour LSPs apply, i.e. it is recommended that the downstream AS of the bypass tunnel and the main LSP are the same AS. Additionally, as two detour LSPs are required for SRLG protection of the upstream link of an egress ASBR and the egress ASBR itself, two bypass tunnels are also required to protect these resources. Note that the LSP-Merge subobject is not used for bypass tunnels as it was the case for detour LSPs because bypass tunnels do not merge with the main LSP at the far-end of the bypass tunnel, but they are terminated at that node.

The difficulty in providing protection with bypass tunnels relies in the selection of appropriate bypasses for the protection of given resources.

To select a bypass tunnel, the PLR has to take a bypass tunnel that it originates and that fulfills the following requirements:

- a) The bypass tunnel must fulfill the appropriate constraints (bandwidth, link affinities, ...).
- b) The bypass tunnel must be disjoint with the link/node/SRLGs to be protected.
- c) The destination of the bypass tunnel must be the next-hop node (resp. next-next-hop node) of the main LSP, or a node further downstream on the path of the main LSP, in case of link protection (resp. node protection).

The first two requirements can be achieved since all required

information is locally available in the PLR. This is because the PLR has established the candidate bypass tunnels, hence it knows the bandwidth and the resources protected by the bypass tunnel. Complying with the third requirement is more difficult. Generally, the PLR must check if the destination of the bypass tunnel belongs to one of the nodes listed in the RRO of the Resv message of the main LSP. Usually the RRO contains interface addresses and the destination of a bypass tunnel may be a different interface address or the node-id of a router. This means that the PLR has to map the addresses listed in the RRO of the main LSP to the destination address of the bypass tunnel. In an intra-area environment this is possible since this information is available in the IGP topology, but in the inter-AS case, this information is not anymore available locally in the PLR. There are multiple methods to solve this problem:

Solution A: use [[NODEID](#)] where the node-id of the routers are put in the RRO of the Resv message of the main LSP and the node-id is also put in the RRO of the Resv message of the bypass tunnel if the destination was not the node-id. In this way, the PLR simply has to compare the node-ids in the RRO of the main LSP with the destination of the bypass tunnel or with the node-id in the RRO of the bypass tunnel.

Solution B: use the interface address that would be recorded in the RRO of the main LSP as destination of the bypass tunnel. For instance, when the link between ASBR1 and ASBR2 is to be protected, the destination address would be the address of the interface on ASBR2 towards ASBR1. If this link is unnumbered, the destination address used is the node-id that is mentioned in the RRO of the main LSP. This is sufficient to identify the common node on the primary and the bypass tunnel. When node protection is to be provided and the destination of the Bypass Tunnel is the next-hop of the protected node (next-next hop from the PLR point of view), the destination of the bypass tunnel should be the address of the interface on the next-next-hop router that goes towards the node being protected. Multiple bypass tunnels must be used in case of parallel links. We also note that a failure of the interface used as destination of the bypass tunnel does not lead to the failure of the bypass tunnel itself (this is in particular important for link protection).

Until now, we supposed that the bypass tunnels were manually configured, with the destination being part of the configuration. But, bypass tunnels can also be signaled automatically when the first main LSP is established. Therefore, we have to determine the destination of these dynamically established bypass tunnels. In case of solution B, the information about the interface addresses in the RRO of the main LSP can be used as a destination address. In case the node-id is put in the RRO, then this node-id can be used.

8. Security Considerations

TBD

Acknowledgments

This work was partially supported by the European Commission within the IST ATRIUM project. The authors would like to thank Dimitri Papadimitriou and Olivier Bonaventure for their useful comments.

References

- [RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", [BCP 9](#), [RFC 2026](#), October 1996.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., Swallow, G., "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), December 2001.
- [ASREQ] Zhang, R., Vasseur, JP., (Editors), "MPLS Inter-AS Traffic Engineering requirements", [draft-ietf-tewg-interas-mpls-te-req-00.txt](#), work in progress.
- [FRR] Pan, P., Atlas, A. (Editors), "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [draft-ietf-mpls-rsvp-lsp-fastreroute-03.txt](#), work in progress.
- [XRO] Lee, CY., Farrel, A., De Cnodder, S., "Exclude Routes - Extension to RSVP-TE", [draft-ietf-ccamp-rsvp-te-exclude-route-00.txt](#), work in progress.
- [HIER] Kompella, K., Rekhter, Y., "LSP Hierarchy with Generalized MPLS TE", [draft-ietf-mpls-lsp-hierarchy-08.txt](#), work in progress.
- [NODEID] Vasseur, J.-P., Ali, Z., Sivabalan, S., "Definition of an RRO node-id subobject", [draft-ietf-mpls-nodeid-subobject-01.txt](#), work in progress.

Authors Addresses


Stefaan De Cnodder
Email: stefaan.de_cnodder@alcatel.be

Cristel Pelsser
Infonet group (FUNDP)
Rue Grandgagnage 21, B-5000 Namur, Belgium
Email: cpe@info.fundp.ac.be

Appendix A: LSP-Merge subobject

The LSP-Merge subobject is a new subobject in the Explicit Route Object (ERO). The procedures defined in [\[RFC3209\] section 5.3.4.1](#) to select the next hop are modified as follows: if after step 3 of the next hop selection process the node finds an LSP-Merge subobject in front of the ERO, i.e. the LSP-Merge subobject is the first subobject in the ERO after removing the subobjects belonging to the local abstract node, then the LSP has to merge with an LSP with the same Session object and LSP ID at the current node, if such an LSP exists. If no such LSP exists, then the detour LSP is rejected and a ResvErr with errorcode TBD is sent to the originating node.

The LSP with which the LSP containing the LSP-Merge subobject merges must be a main LSP, i.e. it may not contain a DETOUR object. In addition the abstract node where the merging occurs must ensure that in case of a failure, the traffic can be switched from the LSP containing the LSP-Merge subobject to a backup LSP that was established by the merging node to protect the main LSP. If these merging conditions cannot be met, the "SRLG protection available" flag inside RRO subobjects, of [appendix B](#), is set to zero. This indicates to the source that SRLG protection is not provided for the main LSP.



L

Set to zero.

Type

TBD.

Length

The length field is set to 4.

Resvd

Set to zero on transmission and ignored on reception.

Appendix B: SRLG protection desired

Currently [[FRR](#)] does not specify how SRLG protection can be requested by the Head-End LSR. One way to do this is to define an "SRLG protection desired" flag in session attribute object. We will not further investigate this since it is outside the scope of this document. i In case two detours or bypass tunnels are available to provide SRLG and node protection, then the "local protection available" flag is set in the corresponding RRO subobject. Similarly, the "bandwidth protection" flag of the RRO subobject is set when both detours or bypass tunnels provide the requested bandwidth. Note that in case of SRLG protection, it is recommended to use sender-template specific detour LSP to avoid merging with other LSPs.