IETF Internet Draft
Proposed Status: Standards Track
Expires: July 2005

Arthi Ayyangar(Editor)
Juniper Networks

Jean-Philippe Vasseur(Editor)
Cisco Systems, Inc.

January 2005

Inter domain GMPLS Traffic Engineering - RSVP-TE extensions

draft-ayyangar-ccamp-inter-domain-rsvp-te-02.txt

Status of this Memo

Copyright Notice

Abstract

This document describes extensions to Generalized Multi-Protocol Label Switching (GMPLS) Resource ReserVation Protocol - Traffic Engineering (RSVP-TE) signaling required to support mechanisms for the establishment and maintenance of GMPLS Traffic Engineering (TE) Label Switched Paths (LSPs), both packet and non-packet, that traverse multiple domains. For the purpose of this document, a domain is considered to be any collection of network elements within a common realm of address space or

path computation responsibility. Examples of such domains include Autonomous Systems, IGP areas and GMPLS overlay networks.

1. Introduction

The requirements for inter-area and inter-AS MPLS Traffic Engineering have been developed by the Traffic Engineering Working Group and have been stated in [INTER-AREA-TE-REQS] and [INTER-AS-TE-REQS] respectively. Many of these requirements also apply to GMPLS networks. The framework for inter-domain GMPLS Traffic Engineering has been provided in [INTER-DOMAIN-FRAMEWORK].

This document presents the RSVP-TE signaling extensions for the setup and maintenance of TE LSPs that span multiple domains. The signaling procedures described in this document are applicable to both packet LSPs ([RSVP-TE]) and non-packet LSPs that use RSVP-TE GMPLS extensions as described in [RSVP-GMPLS]. Three different signaling methods along with the corresponding RSVP-TE extensions and procedures are proposed in this document.

For the purpose of this document, a domain is considered to be any collection of network elements within a common realm of address space or path computation responsibility. Examples of such domains include Autonomous Systems, IGP areas and GMPLS overlay networks ([GMPLS-OVERLAY]).

1.1. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Terminology

ASBR: routers used to connect together ASes of a different or the same Service Provider via one or more Inter-AS links.

Bypass Tunnel: an LSP that is used to protect a set of LSPs passing over a common facility.

ERO: Explicit Route Object

FA: Forwarding Adjacency

FA-LSP: Forwarding Adjacency LSP

LSP: MPLS Label Switched Path

MP: Merge Point. The node where bypass tunnels meet the protected LSP.

NHOP bypass tunnel: Next-Hop Bypass Tunnel. A backup tunnel, which bypasses a single link of the protected LSP.

NNHOP bypass tunnel: Next-Next-Hop Bypass Tunnel. A backup tunnel, which bypasses a single node of the protected LSP.

PLR: Point of Local Repair. The head-end of a bypass tunnel.

Protected LSP: an LSP is said to be protected at a given hop if it has one or multiple associated backup tunnels originating at that hop.

RRO - Record Route Object

TE: Traffic Engineering

TE LSP: Traffic Engineering Label Switched Path

TE-link: Traffic Engineering link

TED: MPLS Traffic Engineering Database

2. Signaling overview

The RSVP-TE signaling of a TE LSP within a single domain is described in [RSVP-TE] and [RSVP-GMPLS]. This document focuses on the RSVP-TE signaling extensions required for inter-domain TE LSP setup and maintenance. Any other extensions that may be needed for routing or path computation are outside the scope of this document.

2.1. Signaling options

There are three ways in which an RSVP-TE LSP could be signaled across multiple domains:

Contiguous - A contiguous TE LSP is a single end-to-end TE LSP that is setup across multiple domains using RSVP-TE signaling procedures described in [RSVP-TE]and [RSVP-GMPLS]. No additional TE LSPs are required to signal a contiguous TE LSP and the same RSVP-TE information for the TE LSP is maintained along the entire LSP path.

Nesting - Nesting one or more TE LSPs into another TE LSP is described in [LSP-HIERARCHY]. This technique can also be used to nest one or more inter-domain TE LSPs into an intra-domain FA-LSP. While

similar to stitching in the control plane, in the data plane, nesting allows for one or more inter-domain LSPs to be transported over a single intra-domain FA-LSP using the label stacking construct.

Stitching - The concept of LSP stitching as well as the required signaling procedures are described in [LSP-STITCHING]. This technique can be used to stitch an inter-domain TE LSP to an intra-domain LSP segment. A inter-domain stitched TE LSP is a TE LSP made up of different TE LSP segments within each domain which are "stitched" together in the data plane so that an end-to-end LSP is achieved in the data plane. In the control plane, however, the different LSP segments are signaled as distinct RSVP sessions which are independent from the RSVP session for the inter-domain LSP.

On receipt of an LSP setup request for an inter-domain TE LSP, the decision of whether to signal the LSP contiguously or whether to nest or stitch it to another TE LSP, depends on the signaled TE LSP characteristics or the local node configuration, when not explicitly signaled. Also, the TE LSP segment or FA-LSP within the domain may either be pre-configured or signaled dynamically based on the arrival of the inter-domain TE LSP setup request.

3. Procedures on the domain boundary node

Whether an inter-domain TE LSP is contiguous, nested or stitched is determined mostly by the signaling method supported by or configured on the intermediate nodes, usually the domain boundary nodes that the inter-domain TE LSP traverses through. It may also depend on certain parameters signaled by the head-end node for the inter-domain TE LSP. When a domain boundary node receives the RSVP Path message for an inter-domain TE LSP setup, it MUST carry out the following procedures before it can forward the Path message to the next hop node,

- apply any locally configured policies
- determine the signaling method to be used based on any desired characteristics signaled by the head-end node of the inter-domain TE LSP or if the signaling method is not explicitly signaled, then determine the signaling method based on local configuration and policies
- depending on the signaling method, carry out any specific ERO procedures, as applicable, as described in the next section
- based on the signaling method to be used, determine the next hop node to forward the RSVP Path message
- in case of nesting or stitching, either find an existing intra-domain TE LSP to carry the inter-domain TE LSP or signal a new one, depending on local policy
- perform any path computations if required. The path computation

procedure itself is outside the scope of this document. The various path computation options are addressed in [INTER-DOMAIN-PATH-COMP]

- in case of any failures (admission control, policy, signaling; etc), originate corresponding error notifications

3.1. Rules on ERO processing

The ERO that a domain boundary node receives in the Path message for an inter-domain TE LSP will be dependent on several factors such as the level of visibility that the head-end node of the inter-domain TE LSP has into other domains, the path computation techniques applied at the head-end node, policy agreements between two domains; etc. Eventually, when the ERO reaches a domain boundary node, the following rules SHOULD be used for ERO processing and signaling. Within a domain, there may be no FA-LSPs or LSP segments. If they are present, then they may originate and terminate on domain boundary nodes. There could also be FA-LSPs and LSP segments that may originate and terminate at other nodes in the domain. In general, these ERO processing rules are also applicable to non-boundary nodes that may participate in signaling the inter-domain TE LSP.

- If there are any policies related to ERO processing for certain LSPs, they SHOULD be applied and corresponding actions should be taken. E.g. if there exists a policy to reject LSP setup request containing ERO with sub-objects identifying nodes within the domain, then a PathErr with the appropriate error code should be sent back
- Section 8.2 of [LSP-HIERARCHY] describes how a node at the edge of a region (domain) processes the ERO in the incoming Path message and uses this ERO, to either find an existing FA-LSP or signal a new FA-LSP using the ERO hops. This also includes adjusting the ERO before sending the Path message to the next hop node. These procedures SHOULD also be followed for nesting or stitching of inter-domain TE LSPs to FA-LSPs or LSP segments respectively. While the domain boundaries are tied to link switching capabilities in [LSP-HIERARCHY], these procedures are also applicable to other domain boundary nodes in the context of this document. E.g. in case of a path computation domain, you have reached the boundary when the ERO hop is no longer reachable via the TE database (TED).
- In case of any failure in processing the ERO hop(s), a Path Error message with appropriate error code ([RSVP-TE]) SHOULD be generated.

3.2. LSP setup failure and crankback

In case of any setup failures along the path due to policy or admission control or other reasons, a corresponding Path Error SHOULD be generated and sent upstream. The propagation of Path Error upstream may be limited to within the domain or it may be sent all the way upstream to the head-end node of the inter-domain TE LSP.

This depends not only on local configuration and ability of a boundary node to do local crankback, but also on any specific parameters requested by the head-end node itself for that LSP. In certain cases, it may be desirable for the head-end node to exert some control on the ability for the boundary nodes to make use of crankback. See [CRANKBACK] for the definition of those bits. When crankback is allowed, the domain boundary node can either decide to forward the Path Error message upstream to the head-end node of the inter-domain TE LSP or try to select another egress boundary node. When crankback is not allowed or if the node has not been configured to do a crankback, then a boundary node, when receiving a Path Error message from a downstream boundary node MUST propagate the Path Error message up to the head-end node of the inter-domain TE LSP.

4. RSVP-TE signaling extensions

The following RSVP-TE signaling extensions are introduced in this document.

4.1. Control of downstream choice of signaling method

In certain mixed environments with different techniques (contiguous, stitched or nested TE LSPs), a head-end node of the inter-domain TE LSP may wish to signal its requirement regarding the signaling method used at the domain boundaries.

[LSP-ATTRIBUTES] defines the format of the Attributes Flags TLV included in the LSP_ATTRIBUTES object carried in an RSVP Path message. The following bit in the Flags TLV is used by the head-end node of the inter-domain TE LSP to restrict the signaling method used by the domain boundary nodes to be contiguous.

0x01 (TBD): Contiguous LSP bit - this flag is set by the head-end node that originates the inter-domain TE LSP if it desires a contiguous end-to-end TE LSP (in the control & data plane). When set, this indicates that a boundary node MUST not perform any stitching or nesting on the TE LSP and the TE LSP MUST be routed as any other TE LSP (it must be contiguous end to end). When this bit is cleared, a boundary node may decide to perform stitching or nesting. A mid-point node not supporting contiguous TE LSP MUST send a Path Error message upstream with an error code of "Routing Problem" and error sub-code=17 (TBD) (Contiguous LSP type not supported). This bit MUST not be modified by any downstream node.

5. Example

5.1. Example topology

In this document, we will consider the following example topology for inter-domain TE LSPs setup and maintenance. Note, however, that inter-domain TE LSP setup across other domains covered by this document will also follow similar signaling procedures. In this example, a domain is an Autonomous system (AS).

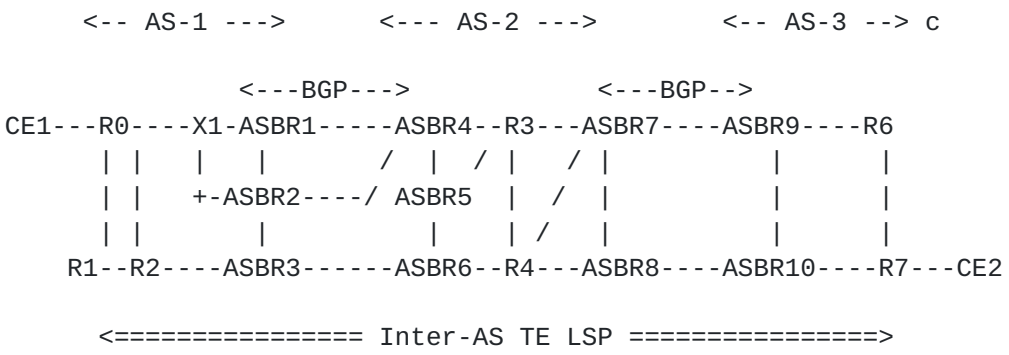

5.1.1. Assumptions

- Three interconnected ASes, respectively AS1, AS2, and AS3. Note that AS3 might be AS1 in some scenarios described in [INTER-AS-TE-REQS].

- The various ASBRs are BGP peers, without any IGP running on the single hop link interconnecting the ASBRs

- Each AS runs an IGP (IS-IS or OSPF) with the required IGP TE extensions (see [OSPF-TE] and [ISIS-TE]). In other words, the ASes are TE enabled. Note that each AS can run a different IGP.

- Each AS can be made of several areas. In this case, the TE LSP will rely on the inter-area TE techniques to compute and set up a TE LSP traversing multiple IGP areas. For the sake of simplicity, each routing domain will be considered as single area in this document, but the solutions described in this document does not prevent the use of multi-area techniques. In fact, these inter-domain solutions are equally applicable to inter-area TE.

- A protected inter-AS TE LSP T1 originated at R0 in AS1 and terminating at R6 in AS3 with following possible paths:

LSP hops: R0-X1-ASBR1-ASBR4-R3-ASBR7-ASBR9-R6

o p1 - a set of loose node hops crossing AS-2

R0-X1-ASBR1(loose)-ASBR4(loose)-ASBR7(loose)-ASBR9(loose)-R6

o p2 - a set of strict interface hops crossing AS-2
R0-X1-ASBR1(loose)-link[ASBR1-ASBR4](strict)-link[ASBR4-R3](strict)
-link[R3-ASBR7](strict)-link[ASBR7-ASBR9](strict)-R6

- A set of backup tunnels:

o B1 from ASBR1 to ASBR4 following the path ASBR1-ASBR2-ASBR4 and protecting against a failure of the ASBR1-ASBR4 link

o B2 from ASBR1 to R3 following the path ASBR1-ASBR2-ASBR3-ASBR6-ASBR5-R3 and protecting against a failure of the ASBR4 node.

o B3 from ASBR1 to ASBR7 following the path ASBR1-ASBR2-ASBR3-ASBR6-ASBR7 and protecting against a failure of the ASBR4 node.

o B4 from R3 to ASBR9 following the path R3-R4-ASBR8-ASBR10-ASBR9 and protecting against a failure of the ASBR7 node.

o B5 from ASBR4 to ASBR9 following the path ASBR4-ASBR8-ASBR10-ASBR9 and protecting against a failure of the ASBR7 node.

5.2. Setup Operation

Let us consider an inter-AS TE LSP setup from R0 to R6, with example paths p1, p2 each. In this example, we will examine the behavior on node ASBR4 which is the boundary node for AS-2, for the different signaling methods.

Contiguous:-

The head-end node, R0, that desires to setup an end-to-end contiguous TE LSP, MAY originate a Path message with LSP_ATTRIBUTES object with the "Contiguous LSP" bit set in the Attributes Flags TLV.

For path p1, additional computation to expand the loose hops may be required at various hops along the LSP path. When the Path message arrives at ASBR4, it may carry out a path computation or use some other means to find the intermediate hops to reach ASBR7. It may then adjust the outgoing ERO and forward the Path message through the intermediate hops in AS-2 to ASBR7.

For path p2, the ERO next hop points to a node within the domain. ASBR4 may then directly forward the Path message to the next hop in the ERO.

Nesting and Stitching:-

When the Path message for the inter-AS TE LSP from R0 to R6, reaches ASBR4, ASBR4 SHOULD first determine from the ERO hops, the boundary node to the domain along the path. In this example, the domain boundary node for all paths is ASBR7. It SHOULD then use the ERO hops upto ASBR7 to find an existing FA-LSP in case of nesting or LSP segment in case of stitching, that satisfies the TE constraints. If there are no existing FA-LSPs or LSP segments and ASBR4 is capable of seting up the FA-LSP or LSP segment on demand, it SHOULD do so using the ERO hops in the Path message of the inter-domain TE LSP. In either case, ASBR4 will adjust the ERO in the inter-domain TE LSP and will forward the Path message directly to the end-point of the FA-LSP or LSP segment using the procedures described in [LSP-HIERARCHY].

In case of path p1, since there are no ERO hops between ASBR4 and ASBR7, and ASBR7 hop is loose, ASBR4 may select any existing FA-LSP (nesting) or LSP segment (stitching) that satisfies the constraints or it may compute a path for the FA-LSP or LSP segment upto ASBR7 or some other intermediate node in AS-2.

In case of path p2, ASBR4 may either select an existing FA-LSP or LSP segment with ERO hops link[ASBR4-R3](strict)-link[R3-ASBR7](strict) or it may compute a new path for the FA-LSP or LSP segment using the above hops. In either case, the ERO hops for the FA-LSP or LSP segment MUST be the same as the signaled strict hops in that domain.

Now, suppose, we have a path p3, as a set of strict node hops crossing AS-2 as defined below,

R0-X1-ASBR1(loose)-ASBR4(strict)-ASBR7(strict)-ASBR9(loose)-R6

In this case, the ERO nexthop at ASBR4 is ASBR7(strict). In this case, ASBR4 will try to find or compute a FA-LSP or LSP segment directly to ASBR7.

The main difference between processing of p1 and p3 for nesting or stitching is that in case of p1, since the ERO nexthop is a loose hop, ASBR4 need not find a FA-LSP or LSP segment directly from ASBR4 to ASBR7. So, there could be multiple FA-LSPs or LSP segments between ASBR4 and ASBR7. On the other hand, for path p3, since ASBR7 is a strict hop, ASBR4 MUST find or signal a FA-LSP or LSP segment that connects ASBR4 and ASBR7.

6. Protection and recovery of inter-domain TE LSPs

6.1. Fast Recovery support using MPLS TE Fast Reroute

[FAST-REROUTE] describes two methods for local protection for a packet TE LSP in case of link, SRLG or node failure. This section describes how these mechanisms work with the proposed signaling solutions for inter-domain TE LSP setup.

6.1.1. Failure within a domain (link or node failure)

The mode of operation of MPLS TE Fast Reroute to protect a contiguous, stitched or nested TE LSP within a domain is identical to the existing procedures described in [FAST-REROUTE]. In case of nested or stitched inter-domain TE LSPs, protecting the intra-domain TE FA-LSP or LSP segment will automatically protect the traffic on the inter-domain TE LSP. No new extensions are required for any of the signaling methods.

6.1.2. Failure of link at domain boundaries

The procedures for doing link protection of the link at domain boundaries is the same for contiguous, nested and stitched TE LSPs.

To protect an inter-domain link with MPLS TE Fast Reroute, a set of backup tunnels must be configured or dynamically computed between the two domain boundary nodes diversely routed from the protected inter-domain link. The region connecting two domains may not be TE enabled. In this case, an implementation will have to support the set up of TE LSP over a non-TE enabled region.

For each protected inter-domain TE LSP traversing the protected link, a NHOP backup must be selected by a PLR (i.e domain exit boundary router), when the TE LSP is first set up. This requires for the PLR to select a bypass tunnel terminating at the NHOP. Finding the NHOP bypass tunnel of an inter-AS LSP can be achieved by analyzing the content of the RRO object received in the RSVP Resv message of both the bypass tunnel and the protected TE LSP(s). As defined in [RSVP-TE], the addresses specified in the RRO IPv4 subobjects can be node-ids and/or interface addresses (with specific recommendation to use the interface address of the outgoing Path messages). The PLR may or may not have sufficient topology information to find where the backup tunnel intersects the protected TE LSP based on the RRO. [NODE-ID] proposes a solution to this issue, defining an additional RRO IPv4 suboject that specifies a node-id address.

Example: The ASBR1-ASBR4 link is protected by the backup tunnel B1 that follows the ASBR1-ASBR2-ASBR4 path

6.1.3. Failure of a boundary node

For each protected inter-domain TE LSP traversing the boundary node to be protected, a NNHOP backup must be selected by the PLR. This requires the PLR to setup a bypass tunnel terminating at the NNHOP. Finding the NNHOP bypass tunnel of an inter-domain TE LSP can be achieved by analyzing the content of the RRO object received in the RSVP Resv message of both the bypass tunnel and the protected TE LSP(s) (see [NODE-ID]). The main difference with node protection, between a protected contiguous inter-domain TE LSP and a protected nested or stitched inter-domain TE LSP is that the PLR and NNHOP (MP) in case of a contiguous TE-LSP could be any node within the domain. However, in case of a nested or stitched TE-LSP the PLR and MP can only be the end-points of the FA-LSP or LSP segment. The consequence is that the backup path is likely to be longer and if bandwidth protection is desired, for instance, ([FAST-REROUTE]) more resources may be reserved in the domain than necessary.

Let us again consider the example topology of section 4.1. The protected inter-domain TE LSP is an inter-AS TE LSP from R0 to R6 with path p1. Also, for nesting or stitching, let us assume that the end-points of the FA-LSP or LSP segment in AS-2 are ASBR4 and ASBR7. This gives rise to the following two scenarios for node protection:

Protecting the boundary node at the entry to a domain :-

Example: protecting against the failure of ASBR4

If the inter-AS TE LSP in this example, is a contiguous LSP, then the PLR is ASBR1 and the NNHOP (MP) could be R3 or any other intermediate node along the LSP path. A backup tunnel B2 may be used to protect the inter-AS TE LSP against failure of ASBR4.

If the inter-AS TE LSP in this example, is nested or sticthed at ASBR4 into an intra-domain TE FA-LSP or LSP segment between ASBR4 and ASBR7, then the PLR is ASBR1 and the NNHOP (MP) is ASBR7. A backup tunnel B3 may be used to protect the inter-AS TE LSP against failure of ASBR4.

Protecting the boundary node at the exit of a domain :-

Example: protecting against failure of ASBR7.

If the inter-AS TE LSP in this example, is a contiguous LSP, then the PLR could be R3 and the NNHOP (MP) is ASBR9. A backup tunnel B4 may be used to protect the inter-AS TE LSP against failure of ASBR7.

If the inter-AS TE LSP in this example, is nested or sticthed at

ASBR4 into an intra-domain TE FA-LSP or LSP segment between ASBR4 and ASBR7, then the PLR is ASBR4 and the NNHOP (MP) is ASBR9. A backup tunnel B5 may be used to protect the inter-AS TE LSP against failure of ASBR7.

6.2. Protection and recovery of GMPLS LSPs

[E2E-RECOVERY] describes the signaling extensions to support end-to-end GMPLS LSP recovery. Signaling methods defined above for inter-domain RSVP-TE LSPs also apply to recovery LSPs signaled for end-to-end protection of inter-domain GMPLS LSPs. Any other protection mechanisms that are developed for GMPLS LSPs SHOULD also take into account inter-domain considerations.

7. Re-optimization of inter-domain TE LSPs

Re-optimization of a TE LSP is the process of moving the LSP from the current path to a more prefered path. This usually involves computation of the new prefered path and make-before-break signaling procedures [RSVP-TE], to minimize traffic disruption. The path computation procedures involved in re-optimization of an inter-domain TE LSP are covered in [INTER-DOMAIN-PATH-COMP].

In the context of an inter-domain TE LSP, since the LSP traverses multiple domains, re-optimization may be required in one or more domains at a time. Again, depending on the nature of the LSP and/or policies and configuration at domain boundaries (or other nodes), one may either always want the head-end node of the inter-domain TE LSP to be notified of any local need for re-optimizations and let the head-end initiate the make-before-break process or one may want to restrict local re-optimizations with the domain.

[LOOSE-REOPT] describes mechanisms that allow,

- The head-end node to trigger on every node whose next hop is a loose hop the re-evaluation of the current path in order to detect a potentially more optimal path. This is done via explicit signaling request: the head-end node sets the "ERO Expansion request" bit of the SESSION-ATTRIBUTE object carried in the RSVP Path message.
- A node whose next hop is a loose-hop to signal to the head-end node that a better path exists. This is performed by sending an RSVP Path Error Notify message (ERROR-CODE = 25), sub-code 6 (Better path exists). This indication may either be sent in response to a query sent by the head-end node or spontaneously by any node having detected a more optimal path.

The above mechanisms SHOULD be used for a contiguous inter-domain TE LSP to allow the head-end node of the inter-domain TE LSP to initiate

make-before-break procedures. For nested or stitched TE LSPs, it is possible to re-optimize the local FA-LSP or LSP segment without involving the head-end node of the inter-domain TE LSP. This will automatically re-route the traffic for the inter-domain TE LSP along the new path, within the domain. Such local re-optimizations, including parameters for re-optimization can be controlled by local policy or configuration in that domain.

8. Security Considerations

When signaling an inter-domain RSVP-TE LSP, an operator may make use of the already defined security features related to RSVP-TE (authentication). This may require some coordination between the domains to share the keys (see RFC 2747 and RFC 3097). Note that this may involve additional synchronization, should the domain boundary LSR be protected with MPLS TE Fast Reroute, since the merge point should also share the key.

For an inter-domain TE LSP, especially when it traverses different administrative or trust domains, the following mechanisms (also see [INTER-AREA-TE-REQS]) SHOULD be provided to an operator :- 1) a way to enforce policies and filters at the domain boundaries to process the incoming LSP setup requests (Path messages) based on certain agreed trust and service levels between domains. 2) a way for the operator to rate limit LSP setup requests or error notifications from a particular domain. 3) a mechanism to allow policy-based outbound RSVP message processing at the domain boundary LSR, which may involve filtering or modification of certain addresses in RSVP objects and messages.

Some examples of the policies described above are:- 1) An operator may choose to implement some kind of ERO filtering policy on the domain boundary LSR to disallow or ignore hops within the domain being identified in the ERO of an incoming Path message. 2) In order to preserve confidentiality, an operator may choose to disallow recording of hops within the domain in the RRO or may choose to filter out certain recorded RRO addresses at the domain boundary LSR. 3) An operator may require the boundary LSR to modify the addresses of certain messages like PathErr or Notify originated from hops within the domain.

Note that the detailed specification of such mechanisms (local implementation) is outside the scope of this document.

9. IANA Considerations

The following values have to be defined by IANA for this document. The registry is, http://www.iana.org/assignments/rsvp-parameters.

9.1. Attribute Flags for LSP_ATTRIBUTES object

The following new flag bit is being defined for the Attributes Flags TLV in the LSP_ATTRIBUTES object. The numeric value should be assigned by IANA.

Contiguous LSP bit - 0x01 (Suggested value)

This flag bit is only to be used in the Attributes Flags TLV on a Path message.

9.2. New Error Codes

The following new error sub-code is being defined under the RSVP error-code "Routing Problem" (24). The numeric error sub-code value should be assigned by IANA.

Contiguous LSP type not supported - sub-code 17 (Suggested value)

This error code is to be used only in a RSVP PathErr.

10. Acknowledgements

The authors would like to acknowledge the input and helpful comments from Adrian Farrel on various aspects discussed in the document.

11. References

11.1. Normative References

[OSPF-TE] Katz, D., Yeung, D., Kompella, K., "Traffic Engineering Extensions to OSPF", RFC 3630 (Updates RFC 2370), September 2003.

[ISIS-TE] Smit, H., Li, T., "IS-IS extensions for Traffic Engineering", RFC 3784

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

[RSVP-GMPLS] L. Berger, et al, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic

Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[LSP-HIERARCHY] Kompella K., Rekhter Y., "LSP Hierarchy with Generalized MPLS TE", (work in progress).

[LSP-STITCHING] Ayyangar A., Vasseur JP., "LSP Stitching with Generalized MPLS TE", (work in progress).

[CRANKBACK] Farrel A. et al, "Crankback Signaling Extensions for MPLS Signaling", (work in progress).

[LSP-ATTRIBUTES] Farrel A. et al, "Encoding of Attributes for Multiprotocol Label Switching (MPLS) Label Switched Path (LSP) Establishment Using RSVP-TE", (work in progress).

[FAST-REROUTE] Ping Pan, et al, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", (work in progress)

[NODE-ID] Vasseur, Ali and Sivabalan, "Definition of an RRO node-id subobject", (work in progress).

[E2E-RECOVERY] J.P. Lang et al, "RSVP-TE Extensions in support of End-to-End GMPLS-based Recovery", (work in progress).

11.2. Informative References

[INTER-AS-TE-REQS] Zhang et al, "MPLS Inter-AS Traffic Engineering requirements", (work in progress).

[INTER-AREA-TE-REQS] LeRoux JL, Vasseur JP, Boyle J. et al, "Requirements for support of Inter-Area MPLS Traffic Engineering", (work in progress).

[INTER-DOMAIN-FRAMEWORK] Farrel A. et al, "A Framework for Inter-Domain MPLS Traffic Engineering", (work in progress).

[INTER-DOMAIN-PATH-COMP] Vasseur JP., Ayyangar A., Zhang R., "Inter-domain MPLS Traffic Engineering LSP path computation methods", (work in progress).

[GMPLS-OVERLAY] G. Swallow et al, "GMPLS RSVP Support for the Overlay Model", (work in progress).

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

[BUNDLING] Kompella K., Rekhter Y., "Link Bundling in MPLS Traffic

Engineering", (work in progress).

[LOOSE-REOPT] Vasseur JP. et al, "Reoptimization of an explicit loosely routed MPLS TE paths", (work in progress).

Author's addresses

Arthi Ayyangar
Juniper Networks, Inc.
1194 N.Mathilda Ave
Sunnyvale, CA 94089
USA
e-mail: arthi@juniper.net

Jean Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MA - 01719
USA
e-mail: jpv@cisco.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.