

Network Working Group
Request for Comments: 4561
Category: Standards Track

J.-P. Vasseur, Ed.
Z. Ali
S. Sivabalan
Cisco Systems, Inc.
June 2006

Definition of a Record Route Object (RRO) Node-Id Sub-Object

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2006).

Abstract

In the context of MPLS TE Fast Reroute, the Merge Point (MP) address is required at the Point of Local Repair (PLR) in order to select a backup tunnel intersecting a fast reroutable Traffic Engineering Label Switched Path (TE LSP) on a downstream Label Switching Router (LSR). However, existing protocol mechanisms are not sufficient to find an MP address in multi-domain routing networks where a domain is defined as an Interior Gateway Protocol (IGP) area or an Autonomous System (AS). Hence, the current MPLS Fast Reroute mechanism cannot be used in order to protect inter-domain TE LSPs from a failure of an Area Border Router (ABR) or Autonomous System Border Router (ASBR). This document specifies the use of existing Record Route Object (RRO) IPv4 and IPv6 sub-objects (with a new flag defined) thus defining the node-id sub-object in order to solve this issue. The MPLS Fast Reroute mechanism mentioned in this document refers to the "Facility backup" MPLS TE Fast Reroute method.

Table of Contents

1. Introduction	2
2. Terminology	4
2.1. Conventions Used in This Document	5
3. Signaling Node-Ids in RROs	5
4. Finding Merge Point	6
5. Security Considerations	7
6. Acknowledgements	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8

1. Introduction

MPLS Fast Reroute (FRR) [FAST-REROUTE] is a fast recovery local protection technique used to protect Traffic Engineering LSPs from link/node/Shared Risk Link Group (SRLG) failure. One or more backup tunnels are pre-established to protect against the failure of a link/node/SRLG. In case of failure, every protected TE LSP traversing the failed resource is rerouted onto the appropriate backup tunnel.

There are several requirements on the backup tunnel path that must be satisfied. First, the backup tunnel must not traverse the element that it protects. In addition, a primary tunnel and its associated backup tunnel should intersect at least at two points (nodes): Point of Local Repair (PLR) and Merge Point (MP). The former is the head-end LSR of the backup tunnel, and the latter is the tail-end LSR of the backup tunnel. The PLR is where FRR is triggered when link/node/SRLG failure happens.

There are different methods for computing paths for backup tunnels at a given PLR. Specifically, a user can statically configure one or more backup tunnels at the PLR with an explicitly configured path, or the PLR can be configured to automatically compute a backup path or to send a path computation request to a PCE (see [PCE-ARCH]).

Consider the following scenario (Figure 1).

Assumptions:

- A multi-area network made of three areas: 0, 1, and 2.
- A fast reroutable TE LSP T1 (TE LSP signaled with the "Local Protection Desired" bit set in the SESSION-ATTRIBUTE object or the FAST-REROUTE object) from R0 to R3.

- A backup tunnel B1 from R1 to R2, not traversing ABR1, and following the R1-ABR3-R2 path.
- The PLR R1 reroutes any protected TE LSP traversing ABR1 onto the backup tunnel B1 in case of ABR1's failure.




Figure 1: Use of Fast Reroute to protect a TE LSP against an ABR failure with MPLS Traffic Engineering Fast Reroute

When T1 is first signaled, the PLR R1 needs to dynamically select an appropriate backup tunnel intersecting T1 on a downstream LSR. However, existing protocol mechanisms are not sufficient to unambiguously find the MP address in a network with inter-domain TE LSP. This document addresses these limitations.

R1 needs to select an existing backup tunnel with the following properties:

1. The backup tunnel intersects with the primary tunnel at the MP. For the sake of illustration, in Figure 1, R1 needs to determine that T1 and B1 intersect at the node R2.
2. The backup tunnel satisfies the primary LSP's request with respect to the bandwidth protection request (i.e., bandwidth guaranteed for the primary tunnel during failure), and the type of protection (link or node failure), as specified in [FAST-REROUTE].

One technique for the PLR to ensure that condition (1) is met consists of examining the Record Route Object (RRO) of the primary tunnel to see whether any of the addresses specified in the RRO correspond to the MP. That said, as per [RSVP-TE], the addresses specified in the RRO IPv4 or IPv6 sub-objects sent in Resv messages can be node-ids and/or interface addresses. Hence, in Figure 1, router R2 may specify interface addresses in the RROs for T1 and B1. Note that these interface addresses are different in this example.

The problem of finding the MP using the interface addresses or node-ids can be easily solved in the case of a single IGP area. Specifically, in the case of a single IGP area, the PLR has the knowledge of all the interfaces attached to the tail-end of the backup tunnel. This information is available in PLR's IGP topology

database. Thus, the PLR can unambiguously determine whether a backup tunnel intersecting a protected TE LSP on a downstream node exists and can also find the MP address regardless of how the addresses carried in the RRO IPv4 or IPv6 sub-objects are specified (i.e., whether using the interface addresses or the node-ids). However, such routing information is not available in the case of inter-domain environments. Hence, unambiguously making sure that condition (1) above is met in the case of inter-domain TE LSPs is not possible with existing mechanisms.

In this document, we define extensions to and describe the use of Resource Reservation Protocol (RSVP) [RSVP, RSVP-TE] to solve the above-mentioned problem. Note that the requirement for the support of the fast recovery technique specified in [FAST-REROUTE] to inter-domain TE LSPs has been specified in [INTER-AREA-TE-REQS] and [INTER-AS-TE-REQS].

2. Terminology

Area Border Routers (ABRs): Border routers used to connect two Interior Gateway Protocol (IGP) areas (areas in OSPF or levels in IS-IS).

Autonomous System Border Router (ASBRs): Border routers used to connect to another AS of a different or the same Service Provider via one or more links inter-connecting between ASes.

Backup Tunnel: The LSP that is used to back up one of the many LSPs in many-to-one backup.

Inter-AS TE LSP: A TE LSP that crosses an AS boundary.

Inter-area TE LSP: A TE LSP that crosses an IGP area.

LSR: Label Switching Router.

LSP: Label Switched Path.

Local Repair: Techniques used to repair TE LSPs quickly when a link, an SRLG, or a node along the TE LSP fails.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

MP: Merge Point. The LSR where one or more backup tunnels rejoin the path of the protected LSP downstream of the potential failure.

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.

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

Reroutable LSP: Any LSP for which the "Local Protection Desired" bit is set in the Flag field of the SESSION_ATTRIBUTE object of its Path messages.

TE LSP: Traffic Engineering Label Switched Path.

2.1. 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].

3. Signaling Node-Ids in RROs

As mentioned above, the limitation that we need to address is the generality of the contents of the RRO IPv4 and IPv6 sub-objects, as defined in [RSVP-TE]. [RSVP-TE] defines the IPv4 and IPv6 RRO sub-objects. Moreover, two additional flags are defined in [FAST-REROUTE]: the "Local Protection Available" and "Local Protection in Use" bits.

In this document, we define the following new flag:

Node-id: 0x20

When set, this indicates that the address specified in the RRO's IPv4 or IPv6 sub-object is a node-id address, which refers to the "Router Address" as defined in [OSPF-TE], or "Traffic Engineering Router ID" as defined in [ISIS-TE]. A node MUST use the same address consistently. Once an address is used in the RRO's IPv4 or IPv6 sub-object, it SHOULD always be used for the lifetime of the TE LSP.

An IPv4 or IPv6 RRO sub-object with the node-id flag set is also called a node-id sub-object. The problem of finding an MP address in a network with inter-domain TE LSP is solved by inserting a node-id sub-object (an RRO "IPv4" and "IPv6" sub-object with the 0x20 flag set) in the RRO object carried in the RSVP Resv message.

An implementation may decide to either:

- 1) Add the node-id sub-object in the RRO carried in an RSVP Resv message and, when required, also add another IPv4/IPv6 sub-object to record interface address.

Example: an inter-domain fast reroutable TE LSP would have the following two sub-objects in the RRO carried in Resv message: a node-id sub-object and a label sub-object. If recording the interface address is required, then an additional IPv4/IPv6 sub-object is added.

or

- 2) Add an IPv4/IPv6 sub-object recording the interface address and, when required, add a node-id sub-object in the RRO.

Example: an inter-domain fast reroutable TE LSP would have the following three sub-objects in the RRO carried in Resv message: an IPv4/IPv6 sub-object recording interface address, a label sub-object, and a node-id sub-object.

Note also that the node-id sub-object may have other applications than Fast Reroute backup tunnel selection. Moreover, it is RECOMMENDED that an LSR recording a node-id address in an IPv4/IPv6 RRO sub-object also set the node-id flag.

4. Finding Merge Point

Two cases should be considered:

- Case 1: If the backup tunnel destination is the MP's node-id, then a PLR can find the MP and suitable backup tunnel by simply comparing the backup tunnel's destination address with the node-id included in the RRO of the primary tunnel.
- Case 2: If the backup tunnel terminates at an address different from the MP's node-id, then a node-id sub-object MUST also be included in the RRO of the backup tunnel. A PLR can find the MP and suitable backup tunnel by simply comparing the node-ids present in the RROs of both the primary and backup tunnels.

It must be noted that although the technique described in this document for selecting an appropriate backup tunnel using the node-id sub-object applies to the case of Inter-area and Inter-AS, in the case of Inter-AS, the assumption is made that the MP's node-id (of the downstream domain) does not overlap with any LSR's node-id present in the PLR's AS.

When both IPv4 node-id and IPv6 node-id sub-objects are present, a PLR may use any or both of them in finding the MP address.

5. Security Considerations

This document does not introduce new security issues. The security considerations pertaining to [RSVP] and [RSVP-TE] remain relevant.

6. Acknowledgements

We would like to acknowledge input and helpful comments from Carol Iturralde, Anca Zamfir, Reshad Rahman, Rob Goguen, and Philip Matthews. A special thanks to Adrian Farrel for his thorough review of this document.

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RSVP] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.
- [RSVP-TE] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [FAST-REROUTE] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.
- [OSPF-TE] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.
- [ISIS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.

7.2. Informative References

- [INTER-AREA-TE-REQS] Le Roux, J.-L., Vasseur, J.-P., and J. Boyle, "Requirements for Inter-Area MPLS Traffic Engineering", RFC 4105, June 2005.
- [INTER-AS-TE-REQS] Zhang, R. and J.-P. Vasseur, "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005.
- [PCE-ARCH] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE) Based Architecture", Work in Progress, April 2006.

Authors' Addresses

J.-P. Vasseur (Editor)
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough , MA - 01719
USA

EMail: jpv@cisco.com

Zafar Ali
Cisco Systems, Inc.
100 South Main St. #200
Ann Arbor, MI 48104
USA

EMail: zali@cisco.com

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario, K2K 3E8
Canada

EMail: msiva@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2006).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

