Network working group
Internet Draft
Intended status: Standards Track
Expires: April 25, 2011

J. Dong
M. Chen
C. Liu
Huawei Technologies
S. Luo
China Telecom

October 25, 2010

Flexible MPLS-TE Fast Reroute
draft-dong-mpls-rsvp-te-plr-designation-01.txt

Abstract

This document defines RSVP-TE extensions which enable the ingress node to designate particular LSRs along the path as Points of Local Repair (PLRs) of the protected LSP, and further indicate the protection type of each PLR. These mechanisms could enhance the control over the establishment of backup LSPs, achieve more flexible TE FRR and also could save the resources needed for establishing and maintaining unnecessary backup LSPs.

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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.

This Internet-Draft will expire on April 25, 2011.

Table of Contents

1. Introduction

Currently the fast reroute mechanisms of RSVP-TE [RFC4090] enable the ingress node of protected LSP to indicate whether local protection is desired and whether node protection is desired for this LSP. However, such indication is relevant to the whole LSP, the ingress node cannot indicate which LSRs on the path are required to be Points of Local Repair (PLRs), and the protection type of each PLR.

This document describes possible solutions for PLR designation in TE fast reroute, and defines simple extensions to RSVP-TE to achieve flexible TE FRR which is backward compatible with RFC 4090.

These mechanisms could provide the operators with more control of the backup LSPs, this is useful when only a subset of the LSRs on the path are required to operate as PLRs. Also, this could avoid unnecessary signaling and bandwidth reservation for protection of components which are not quite likely to fail. Thirdly, this could relieve the burden on LSRs which may not have enough resources to perform local protection functions.

2. 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].

3. Problem Statement

RFC 4090 has defined mechanisms to establish local protection for a particular LSP. The fast reroute mechanisms of RFC 4090 enable the ingress node of the protected LSP to indicate whether local protection is desired and what protection type is needed for this LSP. However, such specification is at the granularity of LSP level, the ingress node cannot explicitly designate which LSRs along the path are required to be PLRs, and the protection type on each PLR.

In some networks some of the links and nodes can be more reliable than the others, e.g. some links may reside in the same building or have redundancy in the physical layer, and some nodes can have good redundancy in both data plane and control plane. Thus there are fewer requirements to protect such links and nodes on LSP level.

Based on the reliability information of the network and service providers' local policy, the operators may prefer to protect only a subset of the links and nodes along the path, thus the ingress node needs to specify particular LSRs as PLRs, and the protection type on each PLR. This can be helpful in many aspects. Firstly, this enables the operators to setup the backup LSPs they need in a more controllable way. Secondly, this could avoid unnecessary signaling and bandwidth reservation for protection of links and nodes which are unlikely to fail. Thirdly, this could relieve the burden on LSRs which may not have enough resources to perform local protection functions.

This document describes possible solutions for PLR designation, and defines simple extensions to RSVP-TE to achieve flexible TE FRR which is backward compatible with RFC 4090.

4. Possible Solutions

4.1. SERO Based Mechanism

GMPLS Segment Recovery [RFC4873] provides one mechanism to specify segment recovery LSPs using SECONDARY_EXPLICIT_ROUTE Object (SERO). An SERO can indicate a recovery LSP's initiator and terminator, standard ERO semantics can optionally be used in SERO to explicitly control the recovery LSP, and a new subobject called Protection is defined to indicate the type of protection or restoration to be provided. Another new Object called SECONDARY_RECORD_ROUTE Object (SRRO) is also defined for this procedure. Detailed mechanisms are specified in section 4 of RFC 4873.

For MPLS networks which support extensions and Objects defined for GMPLS such as SERO, SRRO and PROTECTION, and the operators desire to explicitly specify the path of the recovery LSPs, the SERO based mechanism can be used. Currently there is no detailed specification about the combination use of MPLS-TE FRR [RFC4090] and GMPLS segment recovery [RFC4873]. This section only gives some brief instructions to this mechanism, detailed specification is for further study.

Association between protected LSP and backup LSP: according to RFC 4090, the association is based on the same SESSION Object and the same LSP ID in SENDER_TEMPLATE Object, the only field that varies is the IPv4 (or IPv6) tunnel sender address in SENDER_TEMPLATE Object. The ASSOCIATION Object defined in [RFC4872] MUST not be used.

Designation of PLR and MP: SERO is used to indicate the PLR and the Merge point (MP) of the backup LSP, and optionally to explicitly specify the path of the backup LSP. Note that explicitly designating PLR and MP implies the protection type of TE FRR, i.e. Node Protection or Link Protection.

The PROTECTION subobject in SERO is used to create the PROTECTION object for the recovery LSP. For TE fast reroute, the protection type SHOULD be set to 0x04 (1:N Protection with Extra-Traffic).

Note a node receiving a Path message containing one or more SEROs SHOULD examine each SERO to see if it indicates a local branch point. In scenarios where many backup LSPs are specified using SEROs, this may bring extra burden to nodes which do not have enough control plane resources.

4.2. ERO Based Mechanism

For MPLS networks which do not support the RSVP extensions for GMPLS, the SERO based mechanism may not be applicable. And for operators which do not desire to explicitly specify each node of the backup LSPs, the procedures of SERO based mechanism seems a bit complicated. This section defines simple extensions to Explicit_Route Object (ERO) to achieve flexible PLR designation and protection type indication.

The Explicit_Route Object (ERO) is extended to carry information of PLR designation and type of local protection. The low order bits of the Reserved field in IPv4 prefix and IPv6 prefix subobjects are used as flags to indicate whether the LSR represented by the subobject should operate as a PLR and the desired protection type.

4.2.1. Extensions to IPv4 Prefix Subobject

Two new flags are defined in this subobject. The structure of extended IPv4 prefix subobject is as below:

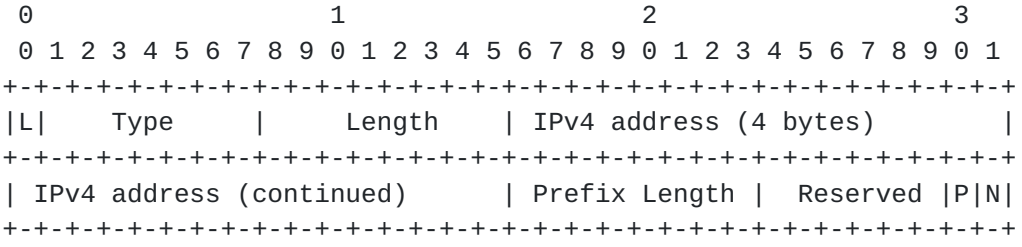

P: Local Protection flag. The P bit indicates whether this subobject is designated as a PLR. It will be set to 0 if the node is designated to be a PLR for the protected LSP, and set to 1 otherwise. If the "Local Protection Desired" flag in the SESSION_ATTRIBUTE Object is not set, no local protection will be used for the whole LSP, and the value of the P bit is insignificant.

N: Node Protection flag. The N bit indicates whether node protection is required for this subobject. It will be set to 1 if node protection is desired, and set to zero if the protection type is indicated by the Node Protection Flag in the SESSION_ATTRIBUTE Object. Note the N bit makes sense only when the "Local Protection Desired" flag in the SESSION_ATTRIBUTE Object is set and the above P bit is set to 0.

4.2.2. Extensions to IPv6 Prefix Subobject

Two new flags are defined in this subobject. The structure of extended IPv6 prefix subobject is as below:

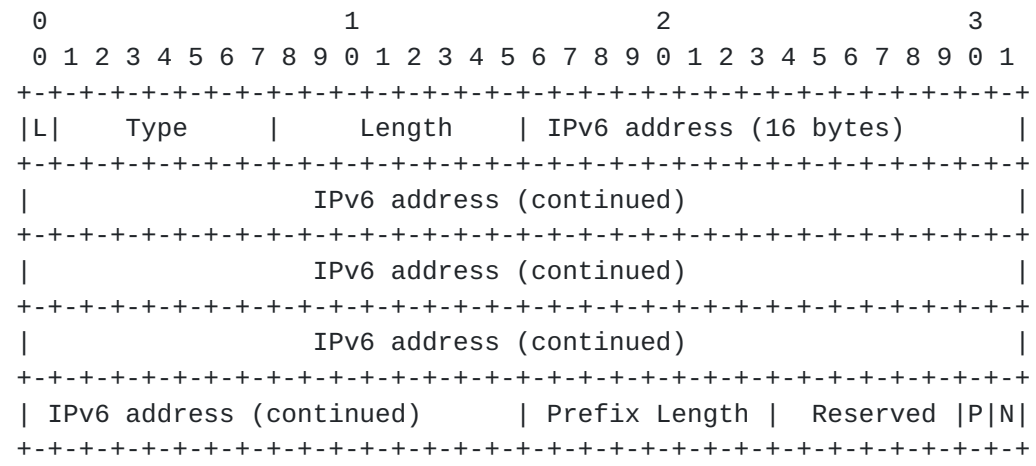

P: Local Protection flag. The P bit indicates whether this subobject is designated as a PLR. It will be set to 0 if the node is designated to be a PLR for the protected LSP, and set to 1 otherwise. If the "Local Protection Desired" flag in the SESSION_ATTRIBUTE Object is not set, no local protection will be used for the whole LSP, and the value of the P bit is insignificant.

N: Node Protection flag. The N bit indicates whether node protection is required for this subobject. It will be set to 1 if node protection is desired, and set to zero if the protection type is indicated by the Node Protection Flag in the SESSION_ATTRIBUTE Object. Note the N bit makes sense only when the "Local Protection Desired" flag in the SESSION_ATTRIBUTE Object is set and the above P bit is set to 0.

4.2.3. Backward Compatibility

The P bit and N bit are designed to be backward compatible with current protection mechanisms. LSRs which do not support this extension will treat these bits as reserved bit and ignore the value of them. When both the 2 bits are set to 0 by head end LSR, the protection behavior of all other LSRs on the path (no matter support this extension or not) is the same as current TE FRR mechanisms.

5. Selection of PLRs and Protection Type

The selection of PLRs and the protection type on each PLR are based on the reliability information of the network and local policy of the service provider. Service providers may have knowledge about which links and nodes in the network are more reliable, and which nodes are not suitable to be PLRs. This kind of information may be obtained by some information advertisement mechanism, or through methods outside the scope of protocols. Based on this information, the operator or the ingress node could designate a subset of LSRs as PLRs and specify the protection type.

6. Operations of ERO Based Mechanism

6.1. Operation of Head End

Based on the result of PLR selection and the required protection type on each PLR, the head-end LSR SHOULD appropriately set the P bit and N bit in corresponding ERO subobjects in the PATH message.

6.2. Operation of Other LSRs

On receipt of a PATH message, the LSR SHOULD check the "Local Protection Desired" and "Node protection desired" flags in the SESSION Attribute Object along with the P bit and N bit in corresponding ERO subobjects, and perform local protection based on these flags.

If some LSR on the path needs to add subobjects to the ERO, it MAY set the P bit and N bit of the subobjects based on local policy.

7. Security Considerations

This document does not introduce new security issues.

8. IANA Considerations

There is no IANA action required by this draft.

9. References

9.1. Normative References

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

[RFC2205] Braden, R., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC3209] 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.

[RFC4090] Pan, P., Swallow, G. and Atlas, A., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC4090, May 2005.

9.2. Informative References

[RFC4872] Lang, J.P., Rekhter, Y. and Papadimitriou, D., "RSVP-TE Extensions in Support of End-to-End GMPLS Recovery", RFC 4872, May 2007.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D. and Farrel, A., "GMPLS Segment Recovery", RFC 4873, May 2007.

Authors' Addresses

Jie Dong
Huawei Technologies Co.,Ltd
Huawei Building, No.3 Xinxi Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

EMail: dongjie_dj@huawei.com

Mach(Guoyi) Chen

Huawei Technologies Co.,Ltd
Huawei Building, No.3 Xinxi Rd.,
Hai-Dian District
Beijing, 100085
P.R. China

EMail: mach@huawei.com

Chun Liu

Huawei Technologies Co.,Ltd
Huawei Building, No.156 Beiqing Rd.
Hai-Dian District
Beijing, 100095
P.R. China

EMail: liuchuner1981@huawei.com

SongFeng Luo

China Telecom
109 West Zhongshan Ave,
Tianhe District, Guanghou, 510630, P.R.C

EMail: luosf08@163.com