Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August 22, 2017

H. Chen
Huawei Technologies
A. Liu
Ciena
T. Saad
Cisco Systems
F. Xu
Verizon
L. Huang
China Mobile
N. So
Tata Communications
February 18, 2017

Extensions to RSVP-TE for LSP Egress Local Protection
draft-ietf-teas-rsvp-egress-protection-07.txt

Abstract

This document describes extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for locally protecting egress nodes of a Traffic Engineered (TE) Label Switched Path (LSP), which is a Point-to-Point (P2P) LSP or a Point-to-Multipoint (P2MP) LSP.

Table of Contents

1. Introduction

RFC 4090 describes two methods for protecting the transit nodes of a P2P LSP: one-to-one and facility protection. RFC 4875 specifies how to use them to protect the transit nodes of a P2MP LSP. However, they do not mention any local protection for an egress of an LSP.

To protect the egresses of an LSP (P2P or P2MP), an existing approach sets up a backup LSP from a backup ingress (or the ingress of the LSP) to the backup egresses, where each egress is paired with a backup egress and protected by the backup egress.

This approach uses more resources and provides slow fault recovery. This document specifies extensions to RSVP-TE for local protection of an egress of an LSP, which overcomes these disadvantages.

1.1. An Example of Egress Local Protection

Figure 1 shows an example of using backup LSPs to locally protect egresses of a primary P2MP LSP from ingress R1 to two egresses: L1 and L2. The primary LSP is represented by star(*) lines and backup LSPs by hyphen(-) lines.

La and Lb are the designated backup egresses for egresses L1 and L2 respectively. To distinguish an egress (e.g., L1) from a backup egress (e.g., La), an egress is called a primary egress if needed.

The backup LSP for protecting L1 is from its upstream node R3 to backup egress La. The one for protecting L2 is from R5 to Lb.

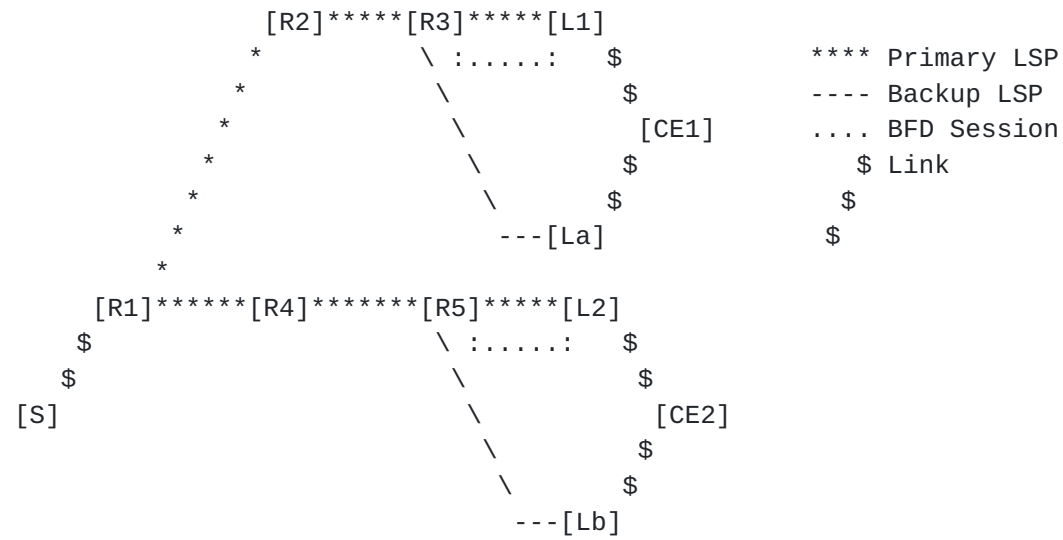

Figure 1: Backup LSP for Locally Protecting Egress

During normal operations, the traffic carried by the P2MP LSP is sent through R3 to L1, which delivers the traffic to its destination CE1. When R3 detects the failure of L1, R3 switches the traffic to the backup LSP to backup egress La, which delivers the traffic to CE1. The time for switching the traffic is within tens of milliseconds.

The failure of a primary egress (e.g., L1 in the figure) may be detected by its upstream node (e.g., R3 in the figure) through a BFD between the upstream node and the egress in MPLS networks. Exactly how the failure is detected is out of scope for this document.

1.2. Egress Local Protection with FRR

Using the egress local protection and the FRR, we can locally protect the egresses, the links and the transit nodes of an LSP. The traffic switchover time is within tens of milliseconds whenever an egress, any of the links and the transit nodes of the LSP fails.

The egress nodes of the LSP can be locally protected via the egress local protection. All the links and the transit nodes of the LSP can be locally protected through using the FRR.

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 RFC 2119.

3. Terminology

This document uses terminologies defined in RFC 2205, RFC 3209, RFC 4090, RFC 4873 and RFC 4875.

4. Protocol Extensions

4.1. Extensions to SERO

The Secondary Explicit Route object (SERO) is defined in RFC 4873. The format of the SERO is re-used.

The SERO used for protecting a primary egress node of a primary LSP may be added into the Path messages for the LSP and sent from the ingress node of the LSP to the upstream node of the egress node. It contains three subobjects.

The first subobject indicates the branch node that is to originate the backup LSP (to a backup egress node). The branch node is the direct upstream node of the primary egress node of the primary LSP if it can provide fast local protection for the primary egress node. The branch node can be a (upstream) node on the primary LSP, but not the direct upstream node if the direct upstream node does not provide any fast local protection against the failure of the primary egress node. In this case, the backup LSP from the branch node to the backup egress node protects against failures on the segment of the primary LSP from the branch node to the primary egress node, including the primary egress node.

The final (third) subobject in the SERO contains the egress node of the backup LSP, i.e., the address of the backup egress node.

The second subobject is a protection subobject defined in RFC 4873, but with some extensions, which include a few of flags for egress local protection and some optional subobjects. These new flags use a few of bits in the existing reserved field. The optional subobjects follow the flags field. The contents of the PROTECTION object is extended in the same way.

The format of the extended protection subobject is defined as follows:

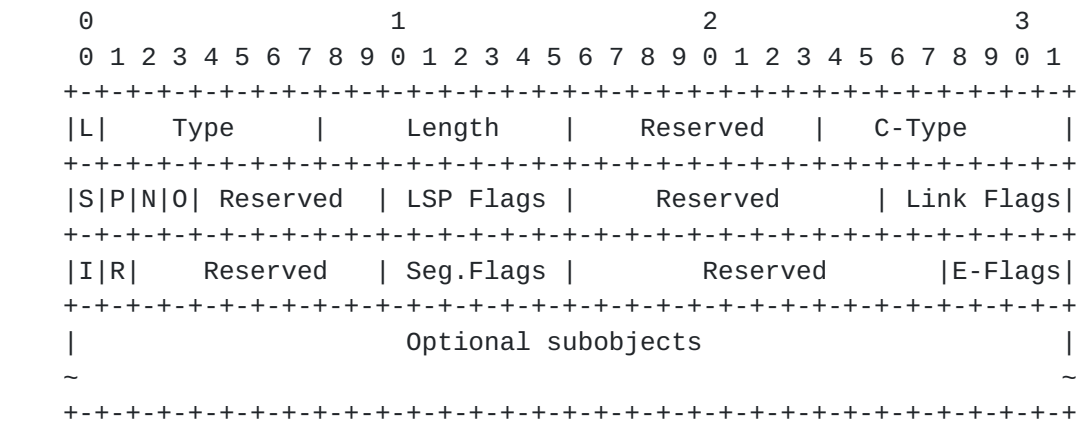

E-Flags are defined for egress local protection.

x01 (Egress local protection bit): It is set (1) to indicate an egress local protection.

x02 (Other sending UA label bit): It is set (1) to indicate another protocol is desired for sending a label as a UA label from a primary egress to a backup egress.

x04 (S2L sub LSP backup desired bit): It is set (1) to indicate S2L Sub LSP (ref to RFC 4875) is desired for protecting an egress of a P2MP LSP.

The other flags are not valid and reset to zero when the egress local protection bit is set.

Five optional subobjects are defined. They are IPv4 and IPv6 primary egress, IPv4 and IPv6 P2P LSP ID, and Opaque Data subobjects. They have the following format:

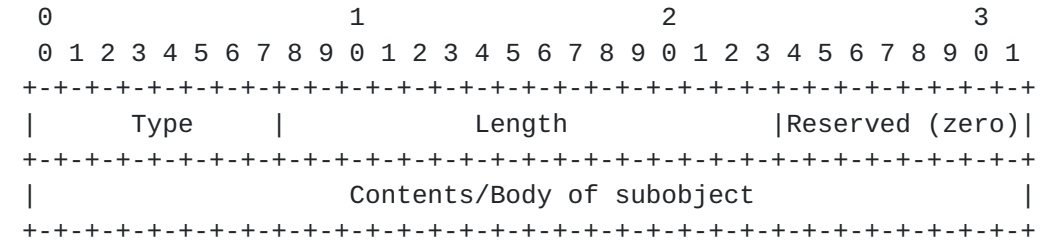

where Type is the type of a subobject, Length is the total size of the subobject in bytes, including Type, Length and Contents fields. The Reserved field MUST be set to zero.

After the upstream node of the primary egress node as the branch node receives the SERO and determines a backup egress node for the primary egress, it computes a path from itself to the backup egress node and sets up a backup LSP along the path for protecting the primary egress node according to the information in the FAST_REROUTE object in the Path message. For example, if facility protection is desired, facility protection is provided for the primary egress node.

The upstream node constructs a PROTECTION object based on the protection subobject in the SERO and adds the object into the Path message for the backup LSP. The object contains a subobject called a primary egress subobject, which indicates the address of the primary egress node.

The upstream node updates the SERO in the Path message for the primary LSP. The protection subobject in the SERO contains a subobject called a P2P LSP ID subobject, which contains the information for identifying the backup LSP. The final subobject in the SERO indicates the address of the backup egress node.

4.1.1. Primary Egress Subobject

There are two primary egress subobjects. One is IPv4 primary egress subobject and the other is IPv6 primary egress subobject.

The Type of an IPv4 primary egress subobject is 1, and the body of the subobject is given below:

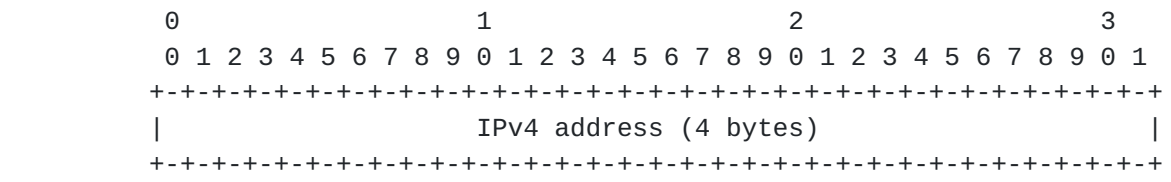

o IPv4 address: IPv4 address of the primary egress node

The Type of an IPv6 primary egress subobject is 2, and the body of the subobject is shown below:

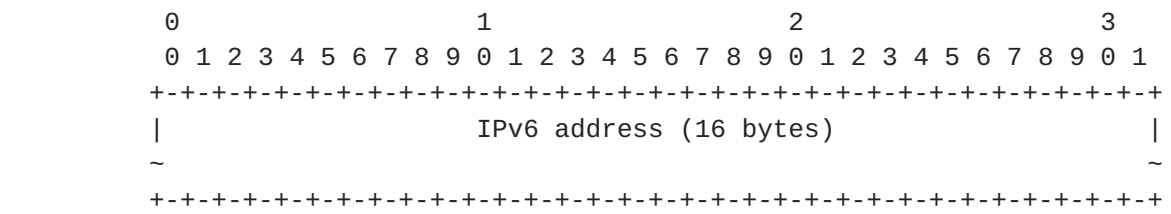

o IPv6 address: The IPv6 address of the primary egress node

4.1.2. P2P LSP ID Subobject

A P2P LSP ID subobject contains the information for identifying a backup point-to-point (P2P) LSP tunnel.

4.1.2.1. IPv4 P2P LSP ID Subobject

The Type of an IPv4 P2P LSP ID subobject is 3, and the body of the subobject is shown below:

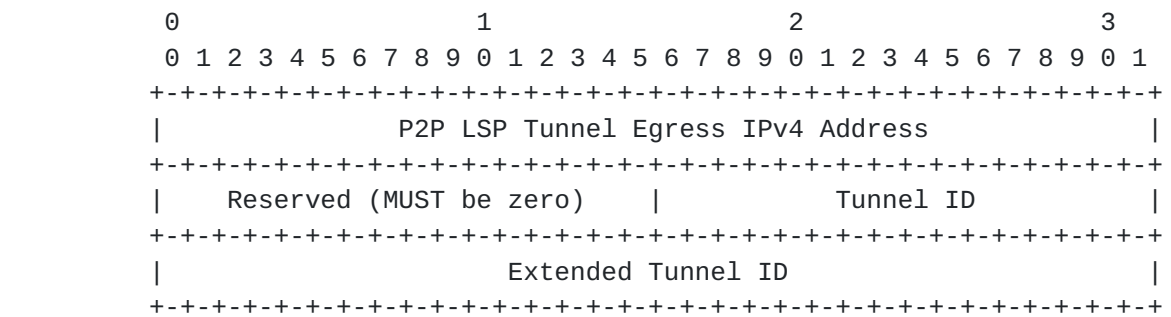

o P2P LSP Tunnel Egress IPv4 Address:
IPv4 address of the egress of the tunnel
o Tunnel ID:
A 16-bit identifier being constant over the life of the tunnel
o Extended Tunnel ID:

A 4-byte identifier being constant over the life of the tunnel

4.1.2.2. IPv6 P2P LSP ID Subobject

The Type of an IPv6 P2P LSP ID subobject is 4, and the body of the subobject is illustrated below:

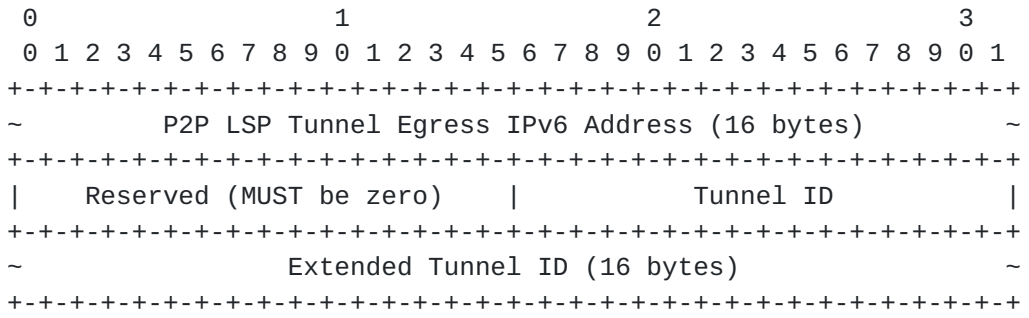

- P2P LSP Tunnel Egress IPv6 Address:
  IPv6 address of the egress of the tunnel
- Tunnel ID:
  A 16-bit identifier being constant over the life of the tunnel
- Extended Tunnel ID:
  A 16-byte identifier being constant over the life of the tunnel

4.1.3. Opaque Data Subobject

A opaque data subobject contains a piece of opaque data. The Type of an opaque data subobject is 5, and the body of the subobject is shown below:

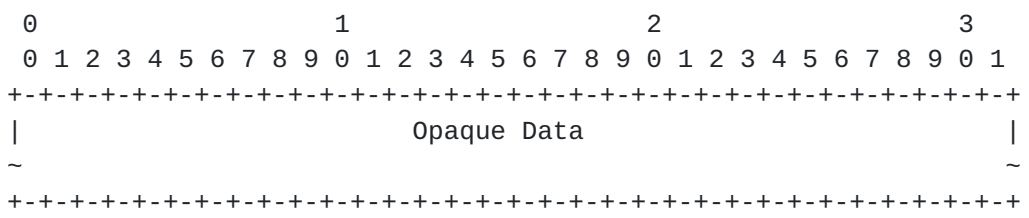

This subobject may be used by a primary egress node (e.g., L1) to send its corresponding backup egress node (e.g., La) the information about forwarding the traffic to a receiver (e.g., CE1) through its upstream node (e.g., R3). It may contain a label as a UA label and a receiver such as CE1. The exact format and meanings of the opaque data are out of scope for this document.

5. Egress Protection Behaviors

5.1. Ingress Behavior

To protect a primary egress of an LSP, the ingress MUST set the "label recording desired" flag and the "node protection desired" flag in the SESSION_ATTRIBUTE object.

If one-to-one backup or facility backup is desired to protect a primary egress of an LSP, the ingress MUST include a FAST_REROUTE object and set the "One-to-One Backup Desired" or "Facility Backup Desired" flag respectively.

If S2L Sub LSP backup is desired to protect a primary egress of a P2MP LSP, the ingress MUST set the "S2L Sub LSP Backup Desired" flag in an SERO object.

If another protocol is desired for sending a label as a upstream assigned label to a backup egress, the ingress MUST set the "Other Sending UA Label" flag.

A backup egress MUST be configured on the ingress of an LSP to protect a primary egress of the LSP if and only if the backup egress is not indicated in another place.

The ingress MUST send a Path message for the LSP with the objects above and the SEROs for protecting egresses of the LSP. For each primary egress of the LSP to be protected, the ingress MUST add an SERO object into the Path message if the backup egress or some options are given. If the backup egress is given, then the final subobject in the SERO containts it; otherwise the address in the final subobject is zero.

5.2. Primary Egress Behavior

To protect a primary egress of an LSP, a backup egress MUST be configured on the primary egress of the LSP to protect the primary egress if and only if the backup egress is not indicated in another place.

If the backup egress is configured on the primary egress of the LSP, the primary egress MUST send its upstream node a Resv message for the LSP with an SERO for protecting the primary egress. It sets the flags in the SERO in the same way as an ingress.

If the LSP carries the service traffic with a service label, the primary egress sends its corresponding backup egress the information about the service label as a UA label and the related forwarding.

5.3. Backup Egress Behavior

When a backup egress node receives a Path message for an LSP, it determines whether the LSP is used for egress local protection through checking the PROTECTION object in the message. If there is a PROTECTION object in the Path message for the LSP and the Egress local protection flag in the object is set to one, the LSP is the backup LSP for egress local protection. The primary egress to be protected is in the primary egress subobject in the PROTECTION object.

When the backup egress receives the information about a UA label and its related forwarding from the primary egress, it uses the backup LSP label as a context label and creates a forwarding entry using the information about the UA label and the related forwarding. This forwarding entry is in a forwarding table for the primary egress node.

When the primary egress node fails, its upstream node switches the traffic from the primary LSP to the backup LSP to the backup egress node, which delivers the traffic to its receiver such as CE using the backup LSP label as a context label to get the forwarding table for the primary egress node and the service label as UA label to find the forwarding entry in the table to forward the traffic to the receiver.

5.4. Transit Node and PLR Behavior

If a transit node of an LSP receives the Path message with the SEROs and it is not an upstream node of any primary egress of the LSP as a branch node, it MUST forward them unchanged.

If the transit node is the upstream node of a primary egress to be protected as a branch node, it determines the backup egress, obtains a path for the backup LSP and sets up the backup LSP along the path. If the upstream node receives the Resv message with an SERO object, it MUST sends its upstream node the Resv message without the object.

The PLR (upstream node of the primary egress as the branch node) MUST extract the backup egress from the respective SERO object in either a Path or a Resv message. If no matching SERO object is found, the PLR tries to find the backup egress, which is not the primary egress but has the same IP address as the destination IP address of the LSP.

Note that if a backup egress is not configured explicitly for protecting a primary egress, the primary egress and the backup egress SHOULD have a same local address configured, and the cost to the local address on the backup egress SHOULD be much bigger than the cost to the local address on the primary egress. Thus primary egress

and backup egress is considered as a virtual node. Note that the backup egress is different from this local address (e.g., from the primary egress' view). In other words, it is identified by an address different from this local address.

After obtaining the backup egress, the PLR computes a backup path from itself to the backup egress and sets up a backup LSP along the path. It excludes the segment including the primary egress to be protected when computing the path. The PLR sends the primary egress a Path message with an SERO for the primary LSP, which indicates the backup egress by the final subobject in the SERO. The PLR puts a PROTECTION object into the Path messages for the backup LSP. The object indicates the primary egress.

The PLR MUST provide one-to-one backup protection for the primary egress if the "One-to-One Backup Desired" flag is set in the message; otherwise, it MUST provide facility backup protection if the "Facility Backup Desired flag" is set.

The PLR MUST set the protection flags in the RRO Sub-object for the primary egress in the Resv message according to the status of the primary egress and the backup LSP protecting the primary egress. For example, it sets the "local protection available" and the "node protection" flag indicating that the primary egress is protected when the backup LSP is up and ready for protecting the primary egress.

5.4.1. Signaling for One-to-One Protection

The behavior of the upstream node of a primary egress of an LSP as a PLR is the same as that of a PLR for one-to-one backup described in RFC 4090 except for that the upstream node as a PLR creates a backup LSP from itself to a backup egress in a session different from the primary LSP.

If the LSP is a P2MP LSP and a primary egress of the LSP is also a transit node (i.e., bud node), the upstream node of the primary egress as a PLR creates a backup LSP from itself to each of the next hops of the primary egress.

When the PLR detects the failure of the primary egress, it switches the packets from the primary LSP to the backup LSP to the backup egress. For the failure of the bud node of a P2MP LSP, the PLR also switches the packets to the backup LSPs to the bud node's next hops, where the packets are merged into the primary LSP.

5.4.2. Signaling for Facility Protection

Except for backup LSP and downstream label, the behavior of the upstream node of the primary egress of a primary LSP as a PLR follows the PLR behavior for facility backup described in RFC 4090.

For a number of primary P2P LSPs going through the same PLR to the same primary egress, the primary egress of these LSPs MAY be protected by one backup LSP from the PLR to the backup egress designated for protecting the primary egress.

The PLR selects or creates a backup LSP from itself to the backup egress. If there is a backup LSP that satisfies the constraints given in the Path message, then this one is selected; otherwise, a new backup LSP to the backup egress is created.

After getting the backup LSP, the PLR associates the backup LSP with a primary LSP for protecting its primary egress. The PLR records that the backup LSP is used to protect the primary LSP against its primary egress failure and MUST include an SERO object in the Path message for the primary LSP. The object MUST contain the backup LSP ID. It indicates that the primary egress MUST send the backup egress the service label as UA label and the information about forwarding the traffic to its destination using the label if there is a service carried by the LSP and the primary LSP label as UA label if the label is not implicit null.

A UA label MAY be sent via another protocol (e.g., BGP). If "Other Sending UA Label" flag is one, the primary egress MUST send the backup egress the UA labels and the information about forwarding the traffic to its destination using the labels through another protocol.

After receiving the Path message with the protection subobject in an SERO, the primary egress MUST include the information about the UA labels in the Resv message with an SERO having a protection subobject containing an opaque data subobject if "Other Sending UA Label" flag is zero. When the PLR receives the Resv message with the information about the UA labels, it MUST include the information in the Path message with a PROTECTION object containing the opaque data subobject for the backup LSP to the backup egress. Thus the UA labels are sent to the backup egress from the primary egress via RSVP.

When the PLR detects the failure of the primary egress, it redirects the packets from the primary LSP into the backup LSP to backup egress and keeps the primary LSP label from the primary egress in the label stack if the label is not implicit null. The backup egress delivers the packets to the same destinations as the primary egress using the backup LSP label as context label and the labels under as UA labels.

5.4.3. Signaling for S2L Sub LSP Protection

The S2L Sub LSP Protection uses a S2L Sub LSP (ref to RFC 4875) as a backup LSP to protect a primary egress of a P2MP LSP. The PLR MUST determine to protect a primary egress of a P2MP LSP via S2L sub LSP protection when it receives a Path message with flag "S2L Sub LSP Backup Desired" set.

The PLR MUST set up the backup S2L sub LSP to the backup egress, create and maintain its state in the same way as of setting up a source to leaf (S2L) sub LSP defined in RFC 4875 from the signaling's point of view. It computes a path for the backup LSP from itself to the backup egress, constructs and sends a Path message along the path, receives and processes a Resv message responding to the Path message.

After receiving the Resv message for the backup LSP, the PLR creates a forwarding entry with an inactive state or flag called inactive forwarding entry. This inactive forwarding entry is not used to forward any data traffic during normal operations.

When the PLR detects the failure of the primary egress, it changes the forwarding entry for the backup LSP to active. Thus, the PLR forwards the traffic to the backup egress through the backup LSP, which sends the traffic to its destination.

5.4.4. PLR Procedures during Local Repair

When the upstream node of a primary egress of an LSP as a PLR detects the failure of the primary egress, it follows the procedures defined in section 6.5 of RFC 4090. It SHOULD notify the ingress about the failure of the primary egress in the same way as a PLR notifies the ingress about the failure of a transit node.

Moreover, the PLR MUST let the upstream part of the primary LSP stay after the primary egress fails through sending Resv message to its upstream node along the primary LSP. The downstream part of the primary LSP from the PLR to the primary egress SHOULD be removed. When a bypass LSP from the PLR to a backup egress protects the primary egress, the PLR MUST NOT send any Path message for the primary LSP through the bypass LSP to the backup egress.

In the local revertive mode, the PLR will re-signal each of the primary LSPs that were routed over the restored resource once it detects that the resource is restored. Every primary LSP successfully re-signaled along the restored resource will be switched back.

6. Considering Application Traffic

This section focuses on the application traffic carried by P2P LSPs. When a primary egress of a P2MP LSP fails, the application traffic carried by the P2MP LSP is delivered to the same destination by the backup egress since the inner label if any for the traffic is a upstream assigned label for every egress of the P2MP LSP.

6.1. A Typical Application

L3VPN is a typical application. An existing solution (refer to Figure 2) for protecting L3VPN traffic against egress failure includes: 1) A multi-hop BFD session between ingress R1 and egress L1 of primary LSP; 2) A backup LSP from ingress R1 to backup egress La; 3) La sends R1 VPN backup label and related information via BGP; 4) R1 has a VRF with two sets of routes: one uses primary LSP and L1 as next hop; the other uses backup LSP and La as next hop.

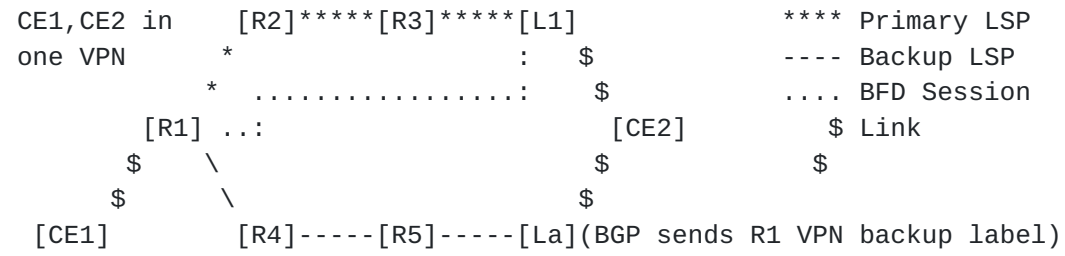

Figure 2: Protect Egress for L3VPN Traffic

In normal operations, R1 sends the traffic from CE1 through primary LSP with VPN label received from L1 as inner label to L1, which delivers the traffic to CE2 using VPN label.

When R1 detects the failure of L1, R1 sends the traffic from CE1 via backup LSP with VPN backup label received from La as inner label to La, which delivers the traffic to CE2 using VPN backup label.

A new solution (refer to Figure 3) with egress local protection for protecting L3VPN traffic includes: 1) A BFD session between R3 and egress L1 of primary LSP; 2) A backup LSP from R3 to backup egress La; 3) L1 sends La VPN label as UA label and related information; 4) L1 and La is virtualized as one. This can be achieved by configuring a same local address on L1 and La, using the address as a destination of the LSP and BGP next hop for VPN traffic.

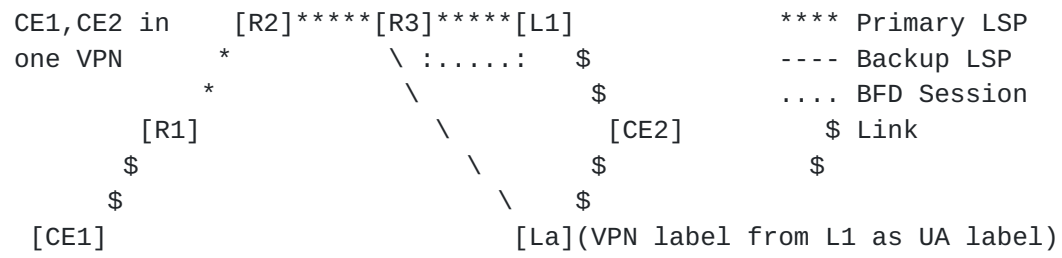

Figure 3: Locally Protect Egress for L3VPN Traffic

When R3 detects L1's failure, R3 sends the traffic from primary LSP via backup LSP to La, which delivers the traffic to CE2 using VPN label as UA label under the backup LSP label as a context label.

6.2. PLR Procedure for Applications

When the PLR gets a backup LSP from itself to a backup egress for protecting a primary egress of a primary LSP, it includes an SERO object in the Path message for the primary LSP. The object contains the ID information of the backup LSP and indicates that the primary egress sends the backup egress the application traffic label (e.g., VPN label) as UA label when needed.

6.3. Egress Procedures for Applications

When a primary egress of an LSP sends the ingress of the LSP a label for an application such as a VPN, it sends the backup egress for protecting the primary egress the label as a UA label. Exactly how the label is sent is out of scope for this document.

When the backup egress receives a UA label from the primary egress, it adds a forwarding entry with the label into the LFIB for the primary egress. When the backup egress receives a packet from the backup LSP, it uses the top label as a context label to find the LFIB for the primary egress and the inner label to deliver the packet to the same destination as the primary egress according to the LFIB.

7. Security Considerations

In principle this document does not introduce new security issues. The security considerations pertaining to RFC 4090, RFC 4875 and other RSVP protocols remain relevant.

Note that protecting a primary egress of a P2P LSP carrying service traffic through a backup egress requires that the backup egress trust the primary egress for the information received for a service label

as UA label.

8. IANA Considerations

IANA is requested to administer the assignment of the new subobject types defined under the PROTECTION object in this document and summarized in this section.

8.1. Definition of PROTECTION Object Reserved Bits

This document defines bits carried in the Reserved field of the PROTECTION object defined in RFC 4873. As no IANA registry for these bits is requested in RFC 4873, no IANA action is required related to this definition.

8.2. New Subobjects

IANA maintains a registry called "Class Names, Class Numbers, and Class Types" under "Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Parameters". IANA is to create and maintain a new registry under PROTECTION object class, Class Number 37, C-Type 2:

o Sub-object type - 37 PROTECTION, C-Type 2

Initial values for the registry are given below. The future assignments are to be made through IETF Review.

Value	Name	Definition
1	IPv4_PRIMARY_EGRESS	Section 4.1.1
2	IPv6_PRIMARY_EGRESS	Section 4.1.1
3	IPv4_P2P_LSP_ID	Section 4.1.2
4	IPv6_P2P_LSP_ID	Section 4.1.2
5	OPAQUE_DATA	Section 4.1.3

9. Co-authors and Contributors

1. Co-authors

Mehmet Toy
Verizon
E-mail: mehmet.toy@verizon.com

Lei Liu
Fujitsu
E-mail: lliu@us.fujitsu.com

Zhenbin Li
Huawei Technologies
Email: lizhenbin@huawei.com

2. Contributors

Boris Zhang
Telus Communications
Email: Boris.Zhang@telus.com

Nan Meng
Huawei Technologies
Email: mengnan@huawei.com

Prejeeth Kaladharan
Huawei Technologies
Email: prejeeth@gmail.com

Vic Liu
China Mobile
Email: liu.cmri@gmail.com

10. Acknowledgement

The authors would like to thank Richard Li, Nobo Akiya, Lou Berger, Jeffrey Zhang, Lizhong Jin, Ravi Torvi, Eric Gray, Olufemi Komolafe, Michael Yue, Daniel King, Rob Rennison, Neil Harrison, Kannan Sampath, Yimin Shen, Ronhazli Adam and Quintin Zhao for their valuable comments and suggestions on this draft.

11. References

11.1. Normative References

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001,

<http://www.rfc-editor.org/info/rfc3209>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <http://www.rfc-editor.org/info/rfc4090>.

[RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <http://www.rfc-editor.org/info/rfc4875>.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and A. Farrel, "GMPLS Segment Recovery", RFC 4873, DOI 10.17487/RFC4873, May 2007, <http://www.rfc-editor.org/info/rfc4873>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

11.2. Informative References

[RFC2205] Braden, R., Ed., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, DOI 10.17487/RFC2205, September 1997, <http://www.rfc-editor.org/info/rfc2205>.

[RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", RFC 5331, DOI 10.17487/RFC5331, August 2008, <http://www.rfc-editor.org/info/rfc5331>.

[RFC4872] Lang, J., Ed., Rekhter, Y., Ed., and D. Papadimitriou, Ed., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, DOI 10.17487/RFC4872, May 2007, <http://www.rfc-editor.org/info/rfc4872>.

[RFC3473] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, DOI 10.17487/RFC3473, January 2003, <http://www.rfc-editor.org/info/rfc3473>.

[FRAMEWK] Shen, Y., Jeyananth, M., Decraene, B., and H. Gredler,

"MPLS Egress Protection Framework",
draft-shen-mpls-egress-protection-framework ,
October 2016.

Authors' Addresses

Huaimo Chen
Huawei Technologies
Boston, MA
USA

Email: huaimo.chen@huawei.com

Autumn Liu
Ciena
USA

Email: hliu@ciena.com

Tarek Saad
Cisco Systems

Email: tsaad@cisco.com

Fengman Xu
Verizon
2400 N. Glenville Dr
Richardson, TX 75082
USA

Email: fengman.xu@verizon.com

Lu Huang
China Mobile
No.32 Xuanwumen West Street, Xicheng District
Beijing, 100053
China

Email: huanglu@chinamobile.com

Ning So
Tata Communications
2613 Fairbourne Cir.
Plano, TX 75082
USA

Email: ningso01@gmail.com