

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 17, 2013

S. Kini, Ed.
S. Narayanan
Ericsson
S. Litkowski
Orange
July 16, 2012

Fast Re-route using extensions to LDP
draft-kini-mpls-frr-ldp-03

Abstract

Label Distribution Protocol (LDP) is widely deployed to signal Label Switched Paths (LSPs) due to its simple operational model. Since LDP establishes LSPs along IGP routed paths, its failure recovery is gated by the interior gateway protocol's (IGP's) re-convergence. Though some mechanisms exist to do Fast Re-route (FRR) of LDP LSPs, they suffer from significant complexity or lack of full coverage or both. This document describes a method to perform FRR of LDP LSPs that retains the simple operational model of LDP. The goal is to provide 100% coverage for all failure scenarios including Shared Risk Link Group (SRLG) failure with recovery characteristics similar to the methods in Resource Reservation Protocol - Traffic Engineering FRR (RSVP-TE-FRR).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on January 17, 2013.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

FRR LDP

July 2012

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
1.1.	Requirements Language	3
2.	Scope	3
3.	Terminology	3
4.	Notation	4
5.	LDP local repair technique	4
6.	Protocol procedures and extensions	9
6.1.	Failure Entity TLV	10
6.1.1.	IP address	11
6.1.2.	SRLG	11
6.2.	Backup Path Vector TLV	11
6.3.	Capability Advertisement TLVs	13
7.	Acknowledgements	14
8.	IANA Considerations	14
9.	Security Considerations	14
10.	References	14
10.1.	Normative References	14
10.2.	Informative References	15
	Authors' Addresses	15

1. Introduction

Label Distribution Protocol (LDP) [[RFC5036](#)] is a widely deployed signaling protocol in MPLS networks due to its simple operational model. It signals LSPs along interior gateway protocol (IGP) routed paths. In case of a failure in the network, the recovery of traffic on LDP LSPs is gated by reconvergence of IGPs. IGPs have relatively slower convergence due to link-state database flooding, route re-computation, route-download etc. An alternative is to use IP Fast Re-route (IPFRR) techniques to provide FRR for LDP LSPs. This includes techniques such as LFA [[RFC5286](#)] that provides an alternate path that can also be used by LDP LSPs. However LFA does not provide full coverage. Other IPFRR methods such as NOT-VIA [[I-D.ietf-rtgwg-ipfrr-notvia-addresses](#)] involve significant complexity. Another approach to protect LDP LSPs is to transport them over RSVP-TE LSPs that are protected using RSVP-TE-FRR [[RFC4090](#)]. This can indirectly protect LDP LSPs. However this has the complexity of deploying an additional protocol RSVP-TE [[RFC3209](#)] and also the network planning to setup the RSVP-TE LSPs that are required to adequately protect the network.

In this document we describe a local-repair mechanism that can provide fast-reroute for LDP LSPs. This mechanism retains the simplicity of LDPs operational model and does not require deployment of additional signaling protocols. It also provides full coverage in all failure scenarios, including shared risk link group (SRLG) failures. The mechanism in this document is henceforth referred to as FRR-LDP. It aims to provide traffic recovery times similar to that provided by RSVP-TE-FRR [[RFC4090](#)]. This mechanism works for a link-state IGP such as OSPF [[RFC2328](#)] or IS-IS[[ISO10589](#)].

1.1. Requirements Language

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](#)].

[2.](#) Scope

This draft is applicable only when per platform label spaces are used. Per interface label spaces are out of scope.

[3.](#) Terminology

Kini, et al.

Expires January 17, 2013

[Page 3]

Internet-Draft

FRR LDP

July 2012

SPT: Shortest Path Tree

PLR: Point of Local Repair. The head-end LSR of a backup-SP LSP.

Backup-SP LSP (BSP LSP): An LDP LSP that provides a backup for a specific failure on the shortest path LDP LSP. The failed entity may be a link, a node or a SRLG. This LSP originates from the PLR.

Backup-SP Merge Point (BSP-MP): The LSR where the Backup-SP LSP is label switched to a label allocated for the shortest path LDP LSP. It need not be downstream of the failed entity.

Exclude-SPT: The shortest path tree from a PLR to a destination, when a particular failure point (link, node, SRLG) is excluded from the topology.

[4.](#) Notation

We use the notation L:X-A to denote the label that LSR A allocates for FEC X for the shortest path LSP to X. We also use the notation Lb:X-A to denote the BSP-LSP label allocated by LSR A for FEC X. The label actions that are setup by LSR A for such a label are explained in detail later in the document.

[5.](#) LDP local repair technique

In a shortest path routed network, unicast traffic to a destination flows along a multi-point to point (MP2P) tree/graph towards that destination. When a failure occurs in such a network, traffic from nodes that are upstream from the failed entity on the MP2P tree gets affected. However, traffic along the tree that is not upstream of the failed entity does not get affected. This is true even when the destination is multi-homed and/or the failed entity consists of several components e.g. SRLG failure.

A PLR protects against a failure detected on its link by re-routing LSPs having that link as a nexthop, along pre-computed and pre-programmed backup paths. The failed entity may be just the link or could be the entire adjacent LSR or even an SRLG with several components. The backup path for a protected LSP must go from the PLR to a LSR (a.k.a Merge Point - MP) that is not upstream of the failed entity in the MP2P tree for that shortest-path protected LSP's destination. Note that the backup path may have to avoid multiple components of the failure e.g., to protect against SRLG failures it must avoid all the components of the SRLG. If there exists a

shortest path LDP LSP that is along a backup path for a LSP, i.e. a shortest-path LDP LSP from the PLR to a LSR that is not upstream of the failure, then it should be used for protection. Such a shortest path LSP is called a Backup Shortest Path LSP (BSP-LSP) and the egress of that LSP which is also the MP is referred to as a backup shortest path merge point (BSP-MP). The PLR must learn the label allocated by the BSP-MP for the protected LSP's destination. The protocol mechanism for to setup the BSP-LSP and for the PLR to learn the label allocated by the BSP-MP for the protected LSP are described in detail later.

To protect against a failure the PLR first swaps the incoming label of the protected shortest-path LDP LSP with the one allocated by the BSP-MP for the destination corresponding to the protected shortest-path LDP LSP. It then tunnels the MPLS frame over the BSP-LSP. All these actions are pre-computed, so that at the time of failure the PLR has to perform a small change of the label forwarding actions to protect the traffic. This helps to recover the traffic in sub 50msec.

A simple example is illustrated below in Figure 1. LSR P protects a LSP to Z against failure of link P-S by using the shortest-path LSP

from P to M i.e. the LSP P-Q-M. LSR P would trigger a protection switch action when link P-S fails that would swap label L:Z-P with L:Z-M, and then tunnel the packet through the shortest-path LSP from P to M by pushing the label L:M-Q and uses nexthop of the shortest-path LSP from P to M i.e. link P-Q. It should be noted that the shortest-path LSP from PLR to BSP-MP i.e., from P to M or P-Q-M that was used as a BSP-LSP would continue to be on the same path even after a failure that it is protecting from. It should also be noted that a BSP-LSP can protect multiple LSPs as long as the BSP-MP is not upstream of the failure entity in the MP2P trees of all the protected LSPs.

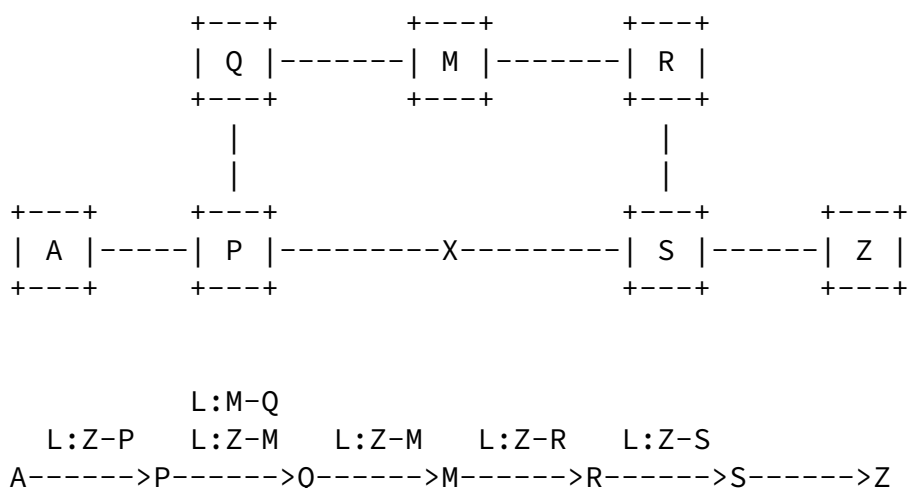
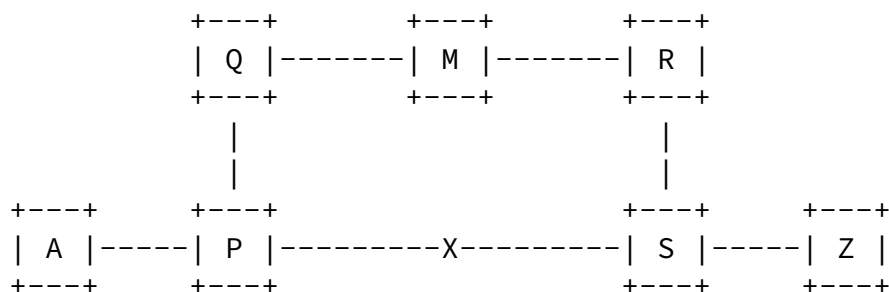


Figure1: Backup Shortest Path LSP

A shortest path LDP LSP may not exist from the PLR to an LSR that is upstream of the failure entity on the MP2P tree. A simple example is illustrated in Figure 2. In such a case the backup path LSP would have non shortest-path forwarding links in it. The backup LSP must take the path P-Q-M that is not a shortest-path LSP from P to M (which is P-S-R-M). However the shortest-path LSP cannot be used since it goes through the failure entity that needs to be protected against i.e. link P-S. The backup path LSP P-Q-M would need label forwarding onto links that are not along the shortest-path LSP from PLR to BSP-MP. In this case LSR Q would have to allocate an additional label to forward the traffic on the backup path i.e. to LSR M. So LSR Q allocates a label Lb:Q-M that has an action of pop when Penultimate Hop Popping (PHP) is the mode of operation and

forward along the link Q-M. The protocol mechanisms to signal label Lb:Q-M from Q to P are described later in this document.



Link Q-M is a high cost link.

```

          Lb:M-Q
    L:Z-P   L:Z-M   L:Z-M   L:Z-R   L:Z-S
A----->P----->Q----->M----->R----->S----->Z

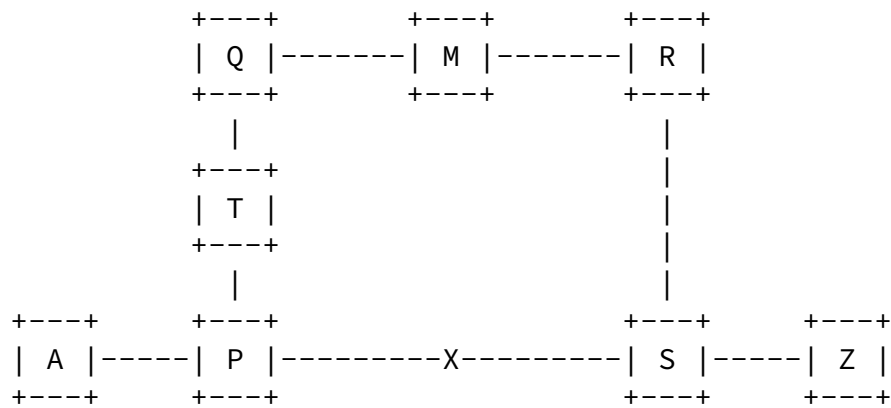
```

Figure 2: Backup Shortest Path LSP with non shortest-path forwarding

By selecting the backup path as the shortest path in the topology with the failure entity removed, the traffic churn after the failure can be minimized. We continue to denote the backup path LSP as BSP-LSP and note that the shortest-path is in the topology with the failure entity removed. If any segment of this path can use existing LDP shortest path LSPs, it would reduce the additional forwarding state needed for the backup LSP. Hence it is recommended that the shortest path LSPs be used as segments in the BSP-LSP wherever possible and keep the non shortest-path links along the BSP-LSP to the minimum. In that case only those LSRs along the backup LSP that routes the backup LSP to a non shortest path link need to allocate additional labels to create the backup LSP.

An example in Figure 3 illustrates how a shortest-path LSP used as a segment in a BSP-LSP reduces forwarding state in intermediate LSRs. LSR P protects the LSP to Z from failure of the link P-S. The BSP-LSP is P-T-Q-M and consists of a shortest-path LSP P-T-Q. Only Q allocates an additional label to forward the packet along the link Q-M. This label has to be advertised to P. LSR T does not have any additional state for the BSP-LSP. For the protected LSP going to Z, P installs an action to swap the label L:Z-P to L:Z-M and push the

label Lb:M-Q in order to tunnel the packet to the BSP-MP i.e. M. LSR P additionally pushes the label L:Q-T with a nexthop of link P-T to follow the shortest-path LSP from P to Q. These actions are triggered on detecting failure of link P-S.

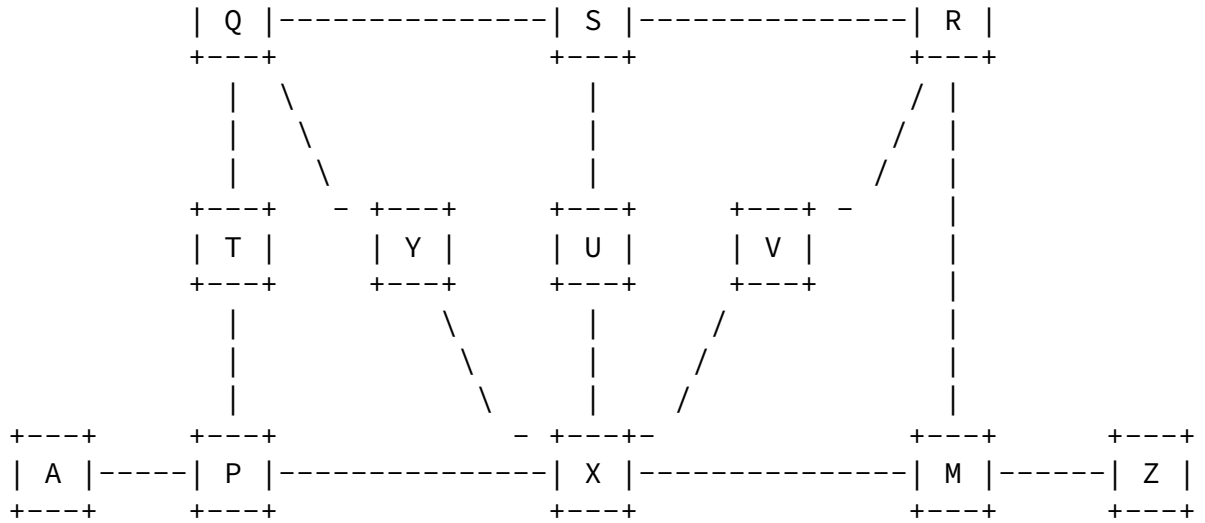


Link Q-M is a high cost link.

L:Q-T
 Lb:M-Q Lb:M-Q
 L:Z-P L:Z-M L:Z-M L:Z-M L:Z-R L:Z-S
 A----->P----->T----->Q----->M----->R----->S----->Z

Figure 3: Backup Shortest Path LSP with non shortest-path forwarding and a shortest-path LSP segment

A slightly more complex topology in Figure 4 illustrates how a combination of shortest path LSPs and non shortest-path links are combined to create a BSP-LSP. PLR P protects against failure of node X by creating a backup LSP P-T-Q-S-R-M. The shortest path LSP Q-S-R can be used as a segment of the BSP-LSP. Since LSRs R and T need to forward the packet on a non shortest-path LSP nexthop, they allocate additional labels to create the BSP-LSP. Other LSRs along the BSP-LSP such as S do not have any additional state for the BSP-LSP.



Link R-M and T-Q are high cost links.

```

          L:R-S
      Lb:M-T  Lb:M-Q  Lb:M-R  Lb:M-R
    L:Z-P  L:Z-M  L:Z-M  L:Z-M  L:Z-M  L:Z-M
A----->P----->T----->Q----->S----->R----->M----->Z

```

Figure 4: Backup Shortest Path LSP with non shortest-path forwarding and a shortest-path LSP segment

There will be a maximum of two additional labels stacked on the packet as it goes along the BSP-LSP. An advantage of using the shortest-path (in the post failure topology) as the backup path is that existing shortest path algorithms can be re-used by simply computing it on topologies with the failure entity removed. Since FRR-LDP is a local repair mechanism, an LSR computes BSP-LSPs only for failures of its immediate nexthops i.e. links/nodes/SRLGs.

It should be noted that the BSP-LSP becomes a single hop LSP when a Loop Free Alternate (LFA) is present. In such cases the LFA could be used and the mechanisms in this draft can be applied only for those cases where an LFA is not available. FRR-LDP can be viewed as an extension to LFA with the additional benefit that it provides full failure coverage against link/node/SRLG failures.

Additional label advertisements are needed for FRR-LDP to function as described above. Firstly the PLR needs the label that the BSP-MP has allocated for the FEC corresponding to the protected LSP. Secondly, labels should be allocated and advertised for the BSP-LSP when an LSR has to route the BSP-LSP along a non shortest-path nexthop. The latter is of course not applicable if the BSP-LSP is itself a

shortest-path LSP, in which case LDP would signal it with existing procedures.

To signal the label of the FEC of the protected LSP from the BSP-MP to the PLR, a targeted LDP session should be used. Note that the PLR is not interested in all label mappings from a BSP-MP. It is interested only in those mappings for which it will tunnel packets to that BSP-MP on a local failure. The PLR should use the Downstream on Demand (DoD) mode to request from the BSP-MP, the specific label mappings that it is interested in.

As explained earlier the BSP-LSP is a LSP that goes from the PLR to the BSP-MP and is the shortest-path from PLR to the BSP-MP in the topology where a failed entity is removed. That failed entity is local to the PLR, though it may have additional components that are non-local to the PLR in case of a SRLG failure. To signal labels for such a LSP using LDP procedures, the label mappings are defined for a combination of the FEC and the failed-entity. The failed entity is identified by a new TLV called "Failure Entity TLV", defined in [Section 6.1](#). It is modeled as a simplified version of the EXCLUDE_ROUTE object XRO [[RFC4874](#)]. The presence of the "Failure Entity TLV" indicates that the LDP message is for a BSP-LSP. To identify the path that the BSP-LSP should take, a new TLV called the "Backup Path Vector TLV" is defined in [Section 6.2](#). This TLV is modeled as a simplified version of the Explicit Route Object (ERO) in RSVP-TE [[RFC3209](#)]. The PLR uses this TLV to route the LSP along a path that is not the shortest-path to the BSP-MP in the current topology but will be the shortest-path in the topology with the failure entity removed. The PLR uses the Label Request message with the FEC of an address of the BSP-MP, the "Failure Entity TLV" and the "Backup Path Vector TLV" to request a label from the downstream LSR. Further details of how the "Backup Path Vector TLV" is processed is described in [Section 6.2](#).

[6](#). Protocol procedures and extensions

To create the BSP-LSP the PLR can compute the path for the BSP-LSP.

If a shortest-path LSP can be used as a BSP-LSP then no extra protocol procedures are required to setup the BSP-LSP. The PLR just needs to get the label corresponding to the FEC of the protected LSP from the BSP-MP. This can be done by setting up a targeted LDP session from PLR to the BSP-MP. This can be a session that is setup dynamically rather than being statically configured. Its setup is initiated by the PLR after the backup path computation that

determines which LSRs becomes BSP-MPs for its local failures. The labels should be advertised from BSP-MP to the PLR in a DoD mode.

If the path calculated for the BSP-LSP is not along a shortest path from the PLR to the BSP-MP, then the PLR initiates a Downstream on Demand (DoD) signaling mechanism defined in LDP [[RFC5036](#)] to setup the BSP-LSP. Some extensions to the DoD mechanism are defined here. A 'Failure Entity TLV' defined in [Section 6.1](#) is included in all messages that signal the BSP-LSP. Together with the FEC it uniquely identifies the LSP that corresponds to the signaling message. The PLR also includes a Backup Path Vector TLV in the label request that indicates the path of the BSP-LSP. The Backup Path Vector TLV is modeled as a simplified version of the ERO defined in RSVP-TE [[RFC3209](#)]. The PLR initiates a Label Request message to the immediate nexthop LSR that in turn generate Label Request messages hop-by-hop by following the Backup Path Vector TLV. When the BSP-MP receives the Label Request it responds with a Label Mapping message. Successive Label Mapping messages are generated by intermediate LSRs till one reaches the PLR and the BSP-LSP is set up. Note that the PLR additionally needs the label from the BSP-MP for the protected LSP's FEC and it continues to get that as explained earlier.

A BSP-LSP along a path that is not a shortest-path from the PLR to the BSP-MP can have segments where it is tunneled through shortest-path LSPs. The PLR signals such segments in the "Backup Path Vector TLV". The head-end LSR of such a segment can decide how the LSP is going to be signaled across the segment. It must support a method where it is signaled on a hop-by-hop basis where the intermediate LSRs on the shortest-path LSP tunnel do not allocate any labels for the BSP-LSP. Alternately the head-end LSR may use a targeted LDP session to the tail-end LSR of the shortest-path segment and send the Label Request message.

[6.1](#). Failure Entity TLV

A Failure entity TLV is defined to specify the topological entity whose failure is being referenced in the signaling messages of the BSP-LSP. Together with the FEC it uniquely identifies a label mapping of the BSP-LSP. This TLV is modeled as a simple version of the EXCLUDE_ROUTE Object defined in Exclude Routes [[RFC4874](#)]. This TLV must be present in the Label Request message, Label Withdraw message, Label Release message, Label Mapping Message and the

Notification message when the BSP-LSP is signaled.

This TLV consists of one of these TLVs

[6.1.1.](#) IP address

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0| Type = IP Prefix (0xTBD) |                               Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IP Address                               |
.                                                                           .
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Prefix Length |   Attribute   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Type - This should be allocated by IANA for IPv4 and IPv6.

Length - This should be 6 for IPv4 and 22 for IPv6

IP Address - This is the IP address.

Prefix Length - This should be length (in bits) of the IP prefix.

Attribute - This should be 0 for link and 1 for the node.

[6.1.2.](#) SRLG

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|0|0| Type = SRLG (0xTBD) |                               Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               SRLG Id                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

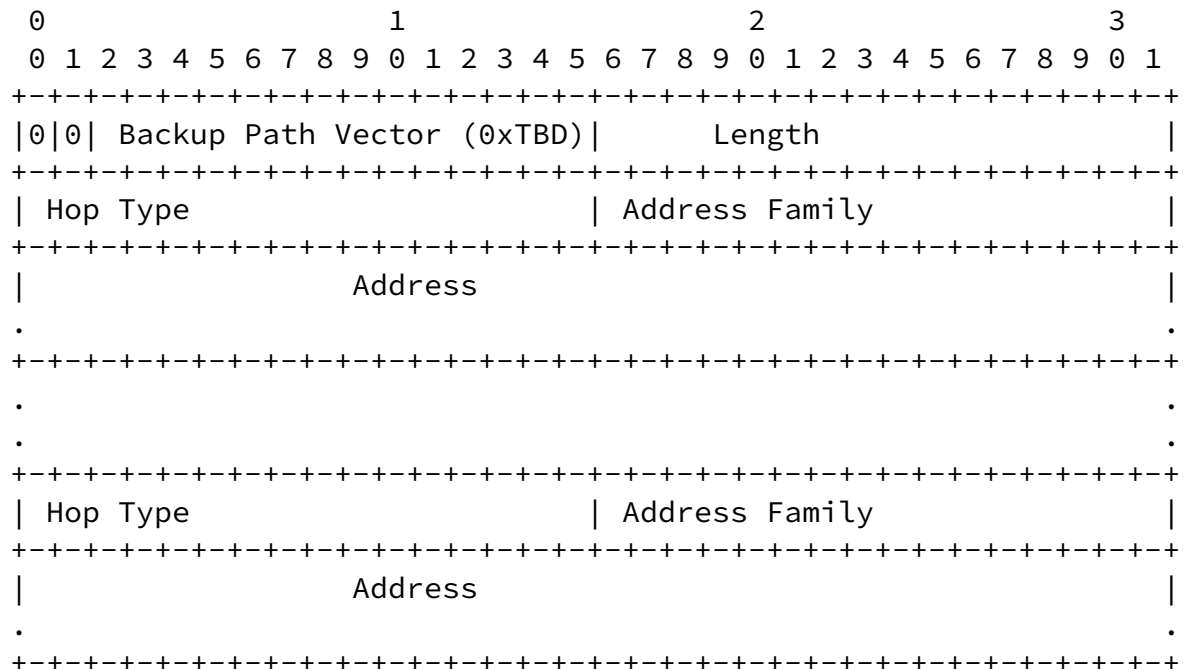
Type - This value should be allocated by IANA.

Length - This should be 4.

SRLG Id - The 32 bit identifier of the SRLG.

6.2. Backup Path Vector TLV

The TLV consists of a list of addresses of LSRs.



Hop Type

0 - Non shortest-path link

1 - Shortest path LSP

2 - Inter-area LSP

Address Family - 0 is for IPv4 and 1 for IPv6

Address - A 4 or 16 octet IP address

This TLV is used in the "Label Request" message. It is processed as follows by examining the first entry in the TLV.

1. If the first entry is its own address and it is the only entry in the Backup Path Vector TLV then it responds to the sender with a Label Mapping message for the FEC corresponding to that address.
2. If the first entry is its own address but it is not the only entry in the Backup Path Vector TLV, then it allocates a label to be sent in response i.e. the Label Mapping message. It also generates a Label Request message towards the IP address in the next entry of the Backup Path Vector TLV. This Label Request will include the Backup Path Vector TLV with the first entry removed. The way the Label Request is generated depends on the "Hop Type"

- A. If the "Hop Type" is 0 (Non shortest-path link) then the Label Request is sent to that directly connected neighboring LSR. When a response is received via a Label Mapping message, a label swap operation from the label it allocated to the label received with forwarding towards that IP address of the neighboring LSR.
- B. If the "Hop Type" is 1 (Shortest path LSP) then the LSR must support generating the Label Request on that shortest-path LSP segment on a hop-by-hop basis. In this method it sends the Label Request to a nexthop LSR along the shortest-path LSP. Alternately it may support sending the Label Request on a targeted LDP session to the tail-end of the shortest-path LSP. When it receives a response to this Label Request message via a Label mapping message, then it installs a label swap operation from the label it allocated to the label received from the downstream LSR with a label forward action to tunnel it over the shortest-path LSP.

- C. If the "Hop Type" is 2 (Inter-area LSP) then this LSR is an area border LSR that must compute a path that avoids the Failure Entity and then generates a Label Request further downstream.
3. If the first entry is not its own address then it generates a Label Request message towards the IP address in the first entry of the Backup Path Vector TLV by choosing one of the LSRs on the nexthop. It does not change the Backup Path Vector TLV in this case. Also it does not allocate a label. When it receives a response in the form of a Label Mapping message it generates a corresponding Label Mapping message to the upstream LSR.

6.3. Capability Advertisement TLVs

A Capability TLV is defined for LDP in accordance to LDP-CAP [RFC5561] to advertise its capability to follow the procedures in this document. This TLV has no additional Capability Data.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
|1|0| Type = Capability (0xTBD) |                               Length |
+-----+-----+-----+-----+-----+-----+-----+-----+
|S| Reserved      |
+-----+-----+-----+-----+-----+-----+-----+

```

To enable PLRs to compute BSP-LSPs the LSRs that are capable of processing the extensions defined in this draft is advertised with

area scope via the link-state IGP. The following extensions are required to advertise this capability:

1. For OSPF an additional bit must be defined in the Router Information Capability bits defined in OSPF-CAP [RFC4970]
2. For ISIS an additional sub TLV to be carried in the CAPABILITY TLV in ISIS-CAP [RFC4971] is defined.

7. Acknowledgements

The authors would like to thank Joel Halpern and TBD for their comments.

8. IANA Considerations

The following LDP TLV Types need assignment.

1. Capability TLV
2. Failure Entity TLV
3. Backup Path Vector TLV

Also a bit for the OSPF Router Information Capability and a TLV value for the ISIS CAPABILITY TLV.

9. Security Considerations

TBD. Security considerations for dynamic LDP sessions.

10. References

10.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), April 1998.
- [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 A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.

- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", [RFC 4874](#), April 2007.
- [RFC4970] Lindem, A., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", [RFC 4970](#), July 2007.
- [RFC4971] Vasseur, JP., Shen, N., and R. Aggarwal, "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", [RFC 4971](#), July 2007.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", [RFC 5286](#), September 2008.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", [RFC 5561](#), July 2009.

[10.2.](#) Informative References

- [I-D.ietf-rtgwg-ipfrr-notvia-addresses]
Bryant, S., Previdi, S., and M. Shand, "IP Fast Reroute Using Not-via Addresses",
[draft-ietf-rtgwg-ipfrr-notvia-addresses-09](#) (work in progress), June 2012.

Authors' Addresses

Sriganesh Kini (editor)
Ericsson

Email: sriganesh.kini@ericsson.com

Srikanth Narayanan
Ericsson

Email: srikanth.narayanan@ericsson.com

Internet-Draft

FRR LDP

July 2012

Stephane Litkowski
Orange

Email: stephane.litkowski@orange-ftgroup.com

