Routing Area Working Group
Internet-Draft
Intended status: Informational
Expires: November 10, 2022

S. Bryant
University of Surrey ICS
U. Chunduri
Intel
T. Eckert
Futurewei Technologies Inc
May 09, 2022

Preferred Path Loop-Free Alternate (pLFA)
draft-bryant-rtgwg-plfa-04

Abstract

Fast re-route (FRR) is a technique that allows productive forwarding to continue in a network after a failure has occurred, but before the network has has time to re-converge. This is achieved by forwarding a packet on an alternate path that will not result in the packet looping. Preferred Path Routing (PPR) provides a method of injecting explicit paths into the routing protocol. The use of PPR to support FRR has a number of advantages. This document describes the advantages of using PPR to provide a loop-free alternate FRR path, and provides a framework for its use in this application.

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 https://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 November 10, 2022.

Table of Contents

1. Introduction

Preferred Path Routing (PPR) [I-D.chunduri-lsr-isis-preferred-path-routing] is a method of introducing explicit paths to a network. Such a path may be any loop-free path between two points in the network that satisfies the

need for which the path was created. The PPR path is not constrained to be the shortest path between any points in the network, although the use of shortest path segments is provided for in order to compress the size of the path description flooded by the routing protocol. The advantages of PPR over alternate methods of creating such paths is described in [I-D.chunduri-lsr-isis-preferred-path-routing].

A packet is carried over the network in an appropriate form using the Preferred Path Routing Identifier (PPR-ID) as the data plane identifier to map the packet to the PPR path, and hence the resources and next-hop (NH). One way of adding a PPR-ID to a packet would be to encapsulate it, but PPR does not restrict the user to the use of encapsulation. How the PPR-ID is carried in the general case is outside the scope of this document. Various methods of adding the PPR-ID to a packet for the purposes of Fast-Reroute (FRR) are described in Section 9.

IP Fast Re-route (IPFRR) Section 1.1 and the methods known at the time of its writing is described in [RFC5714]. A number of later methods are described in [RFC6981], [RFC7490], [RFC7812] and [I-D.ietf-rtgwg-segment-routing-ti-lfa].

This document is a framework describing various methods whereby PPR can be used to provide IP Fast Reroute (IPFRR) paths. PPR can provide IPFRR in a number of ways.

o Signaling pre-computed preferred alternatives for the primary path o Signaling individual segments on the repair path. o Selective overriding of locally computed Loop Free Alternates (LFA) for the NH failure. o Local repair to a Traffic Engineered paths avoiding the need for multi-hop Bidirectional Forwarding Detection (BFD) [RFC5880]. o Micro-loop elimination [RFC5715].

These are described in more detail within this memo.

1.1. A Note on the term IPFRR

The term IP fast re-route (IPFRR) was adopted by the IETF as the general name for best-effort Fast Re-route (FRR) in best effort IP and MPLS networks. This was to distinguish this new work from the then established FRR as described in [RFC4090] which uses RSVP Traffic Engineered (RSVP-TE) MPLS paths [RFC3936].

Within this document the terms IPFRR and FRR are used interchangeably.

2. PPR Overview

PPR works by injecting into the network a path or a graph and a corresponding forwarding identifier (PPR-ID). A node examines each PPR path description and if it is on the path it inserts into the Forwarding Information Database (FIB) an entry for the PPR-ID with the next hop as either the next entry along the PPR path, or if a loose path is specified, the next hop on the shortest path to the next hop along the PPR path. This is described in [I-D.chunduri-lsr-isis-preferred-path-routing].

PPR also has the ability to inject into a network a tree rooted at a node identified a PPR-ID. This is described in [I-D.ce-lsr-ppr-graph]. This graph mechanism provides a compact representation of a set of paths to a given PPR-ID. This works in a similar manner to the linear path case, in which a node on the graph inserts a FIB entry for the PPR-ID with the next hop as either the next node in the graph, or the next hop on the shortest path to the next node in the graph. Clearly the graph needs to be a spanning tree and must not contain a cycle.

In the description of the FRR methods provided in the text, the term encapsulation (and decapsulation) is frequently used in connection with the addition (and removal) of a PPR-ID to be used by the forwarding later to identify to the forwarders the PPR path that the packet needs to traverse to be follow the repair path. Encapsulation is only one of a number of methods that can be used and is used in this memo as a convenience without loss of generality. For more information see Section 9.

3. Preferred Path LFA (pLFA) Deployment Advantages

PPR allows the construction of arbitrary engineered backup paths. In this respect it is like similar to RSVP-TE and Topology Independent Loop-Free Alternates (TI-LFA) [I-D.ietf-rtgwg-segment-routing-ti-lfa]. However, unlike those approaches PPR is applicable to any forwarding plane. For example, it is possible to support MPLS, both IPv4 and IPv6 and Ethernet.

Like Segment Routing (SR) [RFC8402], PPR uses extensions to the existing IGP, however, unlike SR, PPR requires no extension to the data plane. Again, unlike SR, which requires a Segment Identifier (SID) in the network layer header for every non-shortest path forwarding instruction, an arbitrary path does not require expansion of the user data packet beyond that needed for the initial insertion of the PPR-ID. This mitigates the MTU stress that SR introduces to the network.

PPR based IPFRR supports 100% failure coverage similar to RSVP-TE [RFC4090], TI-LFA, Maximally Redundant Trees (MRT) [RFC7812] and Not-Via [RFC6981]. It does not have the coverage restrictions that apply to Loop-Free Alternate (LFA) [RFC5286] and Remote LFA (RLFA) [RFC7490].

Shared Risk Link Groups (SRLGs) make it more difficult find repairs in LFA and RLFA reducing repair coverage. TI-LFA can address this, but only at a cost of expanding the number of SIDs and hence the packet size.

Supporting multiple concurrent failures is difficult in all of the IPRFF approaches except MRT, which can repair two concurrent failures. However unlike MRT, which is constrained by its network wide algorithm, PPR allows individual, arbitrary repair paths to instantiated, for any failure.

In the current TI-LFA design, priority is given to repairing connectivity rather than conforming to the operator traffic policy. A PPR based FFR approach can apply policy to the repaired traffic, including, if required multiple policies to an individual failure.

One of the main advantages of TI-LFA compared to other IPFRR approaches is that it creates repair paths that are congruent with the post convergence path from the Point of Local Repair (PLR) [RFC4090] to the destination. These paths, which may be longer than strictly necessary to reach Q-space [I-D.bryant-ipfrr-tunnels], stop micro-loops from forming along the repair path during re-convergence. PPR can also create these congruent paths without the need to introduce SR into the network.

One of the limitations in TI-LFA, RLFA and LFA is that they do not have a method of selectively creating alternative next-hops or indeed full repair paths based on policy, or traffic engineering information known to the operator. PPR provides a simple way to inject arbitrary paths. It may therefore be used to enhance an existing LFA/RLFA/TI-LFA IPFRR enabled network by selectively injecting paths to provide a repair for business critical links with a policy in the PLR that where provided a PPR path should be preferred over a local calculated LFA based paths.

PPR is applicable to both centralized and PLR computed repair paths each of which has advantages in different circumstances. A centrally computed repair path only requires interaction with one network node which then floods the instruction. This differs from the normal SDN approach which requires interaction with all of the nodes along the path and RSVP-TE which requires interaction with at least one end-point of every repair.

Like TI-LFA, pLFA is based on a small extension to the IGP. It uses the IGP flooding mechanism and in-built state maintenance and consistency checks. This contrasts with RSVP-TE which needs its own separate Signaling and soft-state maintenance method.

The requirement that the pLFA solution addresses is thus the ability to construct repair paths that conform to operator policy without data-plane changes or significant MTU increase, and without introducing any control plane changes other than a small addition to the existing IGP.

A more detailed technical comparison between pLFA and the existing solutions is provided in the technical description of pLFA that follows.

4. Simple Repair Using pLFA

4.1. Link Repair

In this, the most basic, scenario Figure 1 we assume that we have a path A-B-C-D that the packet must traverse. This may be a normal best effort path or a traffic engineered path.

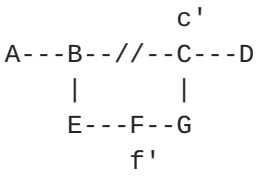

Figure 1: Simple IPFRR Using pLFA

PPR is used to inject the repair path B->E->F->G->C into the network with a PPR-ID of c'. B is monitoring the health of link B->C, for example looking for loss-of-light, or using Bidirectional Forwarding Detection (BFD) [RFC5880]. When B detects a failure it encapsulates the packet to C by adding to the packet the PPR-ID for c' and sending it to E. At C the packet is decapsulated and sent to D. The path C->E->F->G->C may be a traffic engineered path or it may be a best effort path. B may have at its disposal multiple paths to C with different properties for different traffic classes. In this case each path to be used would require its own PPR-ID (c', c'' etc).

In some circumstances, the repair path may be terminated at another point in Q-space or at a node between C and D. For example, in Figure 1 if all costs are 1, F is in Q-space with respect to a B->C failure (F->G->C cost = 2, whilst F->E->B->C cost = 3) and thus the packet can safely be encapsulated and send to F with a PPR-ID of f'. Releasing the packet early in Q-space has two advantages, firstly the

packet can take a shorter path to its destination if one is available rather than traveling to the far side of the failure and then back tracking.

Releasing a packet in Q-space also reduces the size of the PPR path that needs to be advertised, and potentially allows a repair path to be shared among a number of failures. For example in Figure 2 G with PPR-ID g' via B->E->F->G can be used to provide an IPFRR path for the failure of both B->C and B->H.

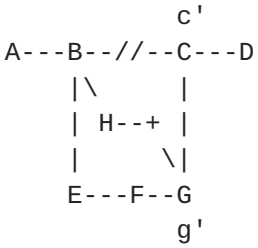

Figure 2: Simple IPFRR Using pLFA With Shared Repair

Shared paths are useful in reducing the number of PPR paths that need to be flooded to support FRR.

Note that where the packet takes the shortest path to the point in Q-space that is closest to the destination, it will be taking a path that is congruent with the post convergence path from the PLR to the destination. This is the path that TI-LFA chooses to avoid its loop-free convergence. However this is not the only loop-free strategy available to a pLFA based solution.

4.2. Node Repair

Consider the network fragment shown in Figure 3 taken from [I-D.bryant-ipfrr-tunnels], and consider that node A needs to deal with the possible failure of node E.

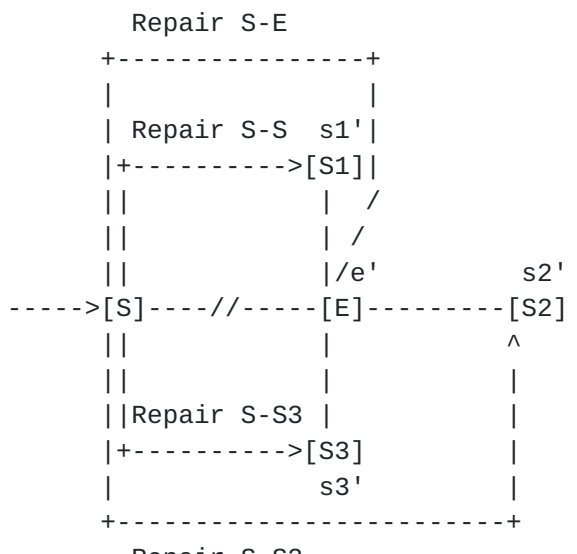

Figure 3: Simple IPFRR - Node Failure

Node S needs the use of four repair paths to address the failure of node E, one repair to each of E's neighbours for which E is on the path to that neighbour. In Figure 3 there are three of these next-next-hop repairs, noted as Repair S-Sx in the figure. In addition a repair to E (Repair S->E) using a path other than along the path S-E> should be installed for traffic to E on the basis that the problem may be a failure of link S->E rather than a failure of the node.

The three repair paths to the next-next-hops of E can be installed as PPR path S-Sx with a PPR-ID of sx'. The link repair for E a PPR path to E which avoids link S-E with a PPR-ID of e'.

4.3. Shared Risk Link Groups

A shared risk link group (SRLG) is a set of links that are believed to have some systematic connection such that when one fails there is a high probability of all of them failing. This occurs, for example, where all of the members of the group run in a common cable duct. Where this relationship is known, and the simultaneous failure does not partition the network,PPR can install paths such that all members of the SRLG are avoided. pLFA has fewer constraints than other methods in constructing arbitrary repair paths in the network. [RFC6981] Section 6.1 describes the SRLG problem as it applies to IPFRR. pLFA can address all of the cases described in [RFC6981].

SRLG avoiding IPFRR paths can be complex. Since a packet can be attracted towards the failure whenever it is released from a strict path, the repair path may need a number of segments to steer it safely into Q-space. If this is done in the data-plane this can

stress the MTU. pLFA creates the path in the control plane and its encapsulation is invariant with respect to the complexity of the path. Furthermore, if the need to reduce the data-plane encapsulation side means that the repair path needs to use a sequence of loose hops it is necessary to determine the behaviour of each router on the chosen path. This contrasts with pLFA which can determine the path using whatever metrics and policy is appropriate, and then simply impose it without any data-plane overhead beyond that needed for a simple repair.

4.4. Local Area Networks

LANs are a special type of SRLG and are solved using the SRLG mechanisms outlined above. With all SRLGs, there is a trade-off between the sophistication of the fault detection and the size of the SRLG. [RFC6981] Section 6.2 describes the LAN problem as it applies to IPFRR. pLFA can address all of the cases described in [RFC6981].

4.5. Multiple Independent Failures

The Multiple Independent Failure cases described in [RFC6981] Section 6.3 will be analyzed in a future version of this document.

4.6. Multi-homed Prefixes

The Multi-Homed Prefix (MHP) problem is described in [RFC5286] Section 6.1, [RFC6981] Section 5.3 and [RFC8518]. MHP will be addressed in a future version of this document.

4.7. ECMP

Equal Cost Multi-Path (ECMP) is a consideration in any IPFRR method that does not use strict paths, and can be both an opportunity and threat. It is an opportunity in that it allows for the repair traffic to be distributed over a number of alternative paths to minimize congestion. If a loose pLFA path is injected into the network, then any available ECMP paths that fulfill the PPR path constraints can be installed following the same procedure used in normal IGP path computation.

However, care must be taken that a packet is not in a position where it is released from a repair at an ECMP point such that one of the ECMP paths is back via the failure. This can never happen if the correct definition of Q-space [RFC7490] is used in calculating the repair path.

5. Repair To A Traffic Engineered Alternate Path

In this approach there are two traffic engineered paths from A to D (Figure 4.

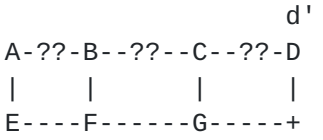

Figure 4: Traffic Engineered IPFRR Using pLFA

The primary path A->B->C->D is protected by a traffic engineered path A->F->G->D (PPR-ID d') with traffic engineered connectors from B (B->F) and C (C->G). The path A->F->G->D and its connectors can be created and injected by any node with access to the IGP, but it is more likely to be created by a traffic engineering controller.

If link B->C fails, B re-routes packets destined for D to the traffic engineered path A->F->G->D via connector B->F. It does this by encapsulating the packet with a PPR-ID of d'.

Clearly there is nothing special needed to get the packet from B to F as they are adjacent but if there is a node say X on the path from B to F an explicit path needs to be created from B to F via X. Normally the repair would be created as a single PPR path (i.e. B->F->G->D) with a PPR-ID of d'. In this approach the repair from A would be A->F->G->D with a PPR-ID of d' also. Similarly C-G-D would again share the PPR-ID d'.

If preferred the repair path could also be constructed using double encapsulation or using an SR approach in which the first segment was B-F with a PPR-ID/SID f' and the second segment was F-D with PPR-ID of d'.

In the example shown in Figure 4 the proposed B-//->C protection path was B->F->G->D. This is node protecting on C since the repair path avoids C. Although link failures tend to be more common than node failures some critical applications would prefer node protection where possible. Node avoidance may not be possible within the network, and may come at a cost of increased path repair path length. However, whether to include node protection and as what cost to accept its inclusion is a matter of network operator policy.

The repair constructed in this section required the inclusion of a set of PPR defined links to construct the repair. PPR has the ability to construct graphs [I-D.ce-lsr-ppr-graph] which can simplify

the specification of the required repair topology. This is discussed in Section 6.

6. Use of a Repair Graph

PPR has the ability to inject graphs into a network as well as linear paths [I-D.ce-lsr-ppr-graph]. PPR graphs specify the paths from a set of nodes to a single node, and are a compact method of representing a set of paths to that destination with shared properties.

6.1. Single Repair Graph

In [I-D.ce-lsr-ppr-graph] the S bit in the PPR Path Description Element (PDE) specifies that a a network node is a Source and a D bit specifies that it is a destination. A graph with all S bits set on the leaves and a D bit on the root is a unidirectional tree.

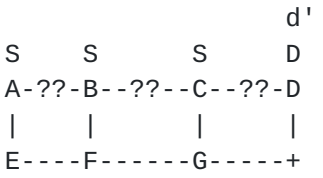

Figure 5: pLFA using PPR Graphs

Consider the network fragment shown in Figure 5. A graph with a PPR-ID d' is constructed attaching each of the nodes A, B, and C to D. Should any of the nodes A, B or C fail the packet can be forwarded on the PPR graph to D with the PPR-ID of d'. In the unidirectional repair graph A, B, and C are all sources (signaled with the S bit set), and D is the only destination (signaled with the D bit set.

6.2. Multiple Disjoint Graphs

Consider Figure 1 from [RFC7812] which illustrates the problem of IPFRR in a network that is 2-connected.

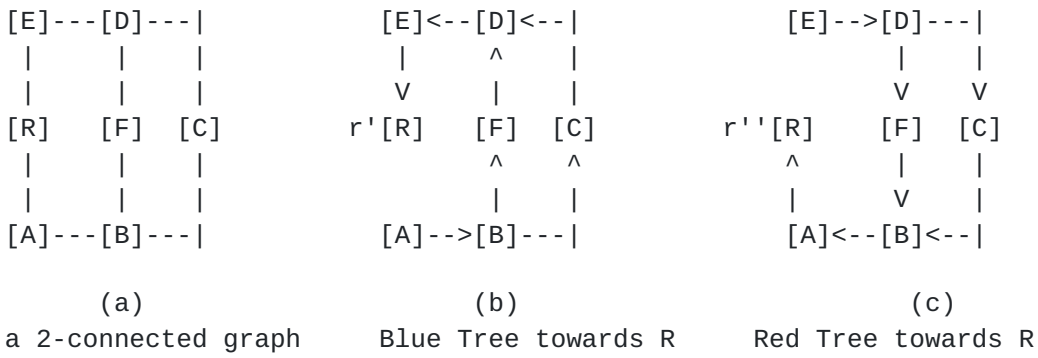

Figure 6: A 2-Connected Network

Figure 6(a) is the full network, and Figure 6(b) and (b) are two corresponding redundant trees from [RFC7812]. Using the Red and Blue trees towards R every node has at least two paths to R. We give R a PPR-ID of r' in the Blue tree and a PPR-ID of r'' in the Red tree. R is the only destination in the PPR graph (D bit set), but all other nodes are sources (S bit set). For clarity this bit setting is not shown in Figure 6.

It is worth noting what happens at nodes B and D in Figure 6(b). B is an ECMP to D via F and C. What happens at node B is a matter of implementation and operator preference. Either B can choose one of the next-hops, or it use them as an ECMP pair. It can also use the availability of the pair to protect against B->F or B->C being an unexpected SRLG with respect to link A->R. D is a merge point for traffic destined for r' arriving from F and from C. It simply forwards the traffic to r' as normal. Similarly in Figure 6(c) D can sent traffic to r'' via F or C.

Whilst in this example the Red and Blue trees use exactly the same links and nodes used by the main topology, a repair graph could use available nodes and links outside this network fragment.

Now consider Figure 2 from [RFC7812] which illustrates the problem of IPFRR in a network that is not 2-connected.

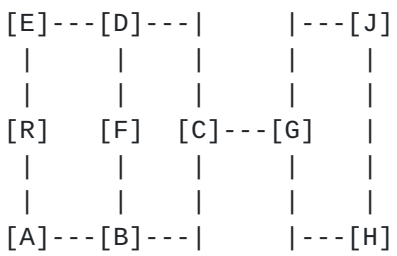

(a) a graph that is not 2-connected

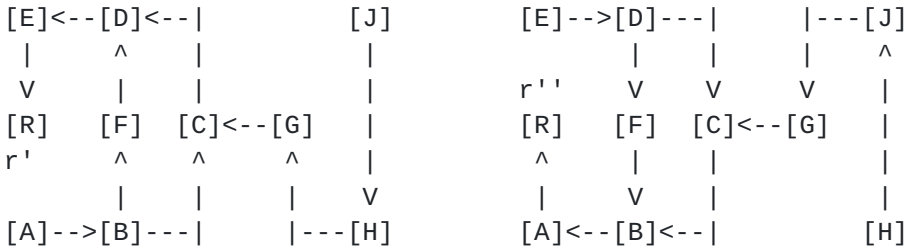

(b) Blue Tree towards R (c) Red Tree towards R

Figure 7: A Network That Is Not 2-Connected

Again there are two paths (with PPR-IDs r' and r'') to R from all nodes except that G, J and H all depend on link G->C and node C which is a single point of failure in the network.

Again note that B in the Blue tree and D in the Red tree has two paths to r' and r'' respectively that it may use according to configuration or preference.

7. Centralized and Decentralized Approaches

pLFA paths can be established through both centralized and decentralized approaches.

A centralized system has a more holistic view of the network and its policies, its resource constraints and resource usage. A decentralized system is inherently more resilient to failure and is a good fit where the network is a simple best effort network as is commonly deployed.

A centralized system gathers the network state, just as any SDN system does, and computes the FRR paths needed. However, unlike normal SDN operation where the controller needs to individually instruct every entity on the path for every path, in a PPR network it is only necessary to inject the PPR path at one point. In practice, for reliability, it would inject the PPR paths in a small number of places, and the naturally reliability of the IGP would ensure the

complete distribution of the paths. Furthermore, the system collecting the network state would naturally send the PPR LSPs back to the SDN controller providing quality assurance that the FRR paths had been distributed.

In a decentralized approach the pLFA path is computed within the network, normally by the PLR. Further details of this approach will be provided in a future version of this document.

8. Independence of operation

Each PPR path is independent of all other paths in the network. This means that there is no constraint on how the path is calculated, and a different algorithm can be used on every path. Some of the other FRR approaches have this property, but not all. For example LFA is constrained by the properties of the base IGP as to a large degree is RLFA. PPR can incorporate best effort segments if required, but from a data-plane perspective there is no advantage in doing so. In this case there is a dependence on the path choice in the base routing protocol.

MRT and Not-Via can use any algorithm to calculate the repair, but it needs to be common across the network, although the expectation in the case of Not-Via is that the algorithm would be a Dijkstra based SPF calculation. In both these cases to change the algorithm would require turning off FRR for the whole network, re-configuring and then restarting FRR.

RSVP-TE based FRR can specify any path, but at the cost of maintaining the soft-state.

A PLR in a TI-LFA or any SR based approach can also compute paths independent of each other, but they tend to need to do this as a concatenation of a series of shortest paths in order to reduce the number of SIDs they need to form the path. TI-LFA is thus highly dependent on the underlying best effort paths.

pLFA can be used as a method of converting classic LFA or RLFA to full coverage by providing the paths that these methods are unable to support, or to provide any the sub-paths needed to reduce the number of TI-LFA SIDs.

9. Data-plane Considerations

This section is a survey of a number of data-planes in each case considering how a PPR-ID could added to map the packet to required FRR path.

9.1. Traditional IP

Where the data-plane is "traditional" IP the user packet needs to be encapsulated such that the outer IP address is the PPR-ID. Any preferred encapsulation can be used such as: IP in IP, IP in GRE, or IP in UDP.

The tunnel capabilities of a node can be advertised using the method described in [I-D.xu-isis-encapsulation-cap] allowing different tunnel types to be used for different PPR paths, depending on the capability of the various nodes in the network.

A common operational issue with this type of encapsulation for IPFRR has been the shortage of IP addresses. However this is not an issue in an IPv6 network.

9.2. Segment Routing over an IPv6 Data Plane (SRv6)

Where the data-plane is SRv6 [RFC8754] pLFA would be used to steer a packet towards the next segment end-point. Clearly an extra level of IP encapsulation could be used Section 9.1, but that expands the packet by adding at least 36 octets.

Where the packet is a "traditional" IP packet, and the repair end-point is SRv6 capable, an alternative to the methods described in Section 9.1 is to insert an SRH into the IP packet setting the SID in the SRH to the original packet DA and replacing the outer DA with the PPR-ID. If this method is used the semantics of the PPR-ID must include the reconstruction of the packet, by replacing the DA with the original DA retrieved from the SRH and the removal of the SRH.

9.3. MPLS

Where the data-plane is MPLS any encapsulation needed is tiny (a label push), but the exact action depends on the repair strategy, and there is the usual FRR problem of the setting of the new value for the top label prior to pushing the PPR-ID label.

Where the FRR path terminates at an MPLS node other than the network egress provider edge (PE) in the type of pLFA repair described in Section 4, the original top label needed to be set to the label the node was expecting.

Consider the network fragment shown in Figure 1. This is straight forward case because node B swaps the top label the label it would have used without the failure and then pushes the label that corresponds to c'. If the repair strategy had been to exit Q-space at the earliest opportunity for example at F, then B would have

needed to know what label F required to reach the destination. A very similar problem occurs when a node repair is undertaken Figure 3, where S needs to know the label that the next-next-hops (S1, S2 and S3) need to reach the destination.

Where the traffic is being moved to a new path terminating at the egress PE as shown in Figure 4, the problem much simpler and only requires the swapping of the top label with the label that represents d'.

10. Loop Free Convergence

Whilst IPFRR puts in place a temporary network repair, eventually the network needs to re-converge around the surviving network components. During this phase there is a danger that micro-loops will form and disrupt the traffic flowing across the network. A similar problem can occur when the failed component returns to service, or when a new component is introduced into the network. [RFC5715] describes the problem of loop-free convergence in detail and examines the methods known at the time of its writing. Since that time [RFC8333] has proposed a timer based loop mitigation (but not elimination) process, and [I-D.ietf-rtgwg-segment-routing-ti-lfa] has proposed that by making the IPFRR path congruent with the post convergence path loops can be eliminated along the repair path. However whilst these mitigation techniques address component failure, neither are targeted at the repair/new component case.

These problems only effect best effort paths and path segments, fully defined paths do not have this problem.

A network using pLFA is compatible with all of the know loop-free convergence and loop mitigation approaches.

11. OAM Considerations

PPR may also be used in a way that provides an alternative to running multi-hop BFD from ingress on a traffic engineered (TE) path with reducing the complexities that arise from echo reply false alarms. In this use case pLFA works by locally detecting the failure and transferring the traffic to preferred TE backups which are in time replaced by the newly computed TE paths to the same PPR-ID.

12. Privacy Considerations

As noted in Section 13 pLFA paths are constrained by the routing domain and thus the traffic will be no more subject to observation than it would in normal operation. Indeed PPR has the capability to constrain the path of the traffic more tightly than other IPFRR

approaches. pLFA therefore does not reduce the privacy of user traffic on the network.

13. Security Considerations

The security considerations of PPR are discussed in [I-D.chunduri-lsr-isis-preferred-path-routing] which in turn refers the reader to the security considerations of the underlying routing protocol and the data-plane in use. The pLFA application of PPR to IPFRR introduces no additional security regarding PPR itself.

General IPFRR security considerations are discussed in [RFC5714] and these apply to this solution.

One further consideration, is the whether policy that applied to the original path needs to be applied to the repair path. The decision is operator and application specific, however pLFA is better than some other IPFRR solution in that it is possible to precisely choose the repair path.

IPFRR is deployed within the scope of the routing protocol that underpins it which limits the security vulnerability. Furthermore it is unlikely that IPFRR would be deployed outside a well managed network. These restrictions in-turn significantly mitigate any security threat.

14. IANA Considerations

This document makes no IANA requests.

15. References

15.1. Normative References

[I-D.ce-lsr-ppr-graph]
Chunduri, U. and T. Eckert, "Preferred Path Route Graph Structure", draft-ce-lsr-ppr-graph-04 (work in progress), September 2020.

[I-D.chunduri-lsr-isis-preferred-path-routing]
Chunduri, U., Li, R., White, R., Contreras, L. M., Tantsura, J., and Y. Qu, "Preferred Path Routing (PPR) in IS-IS", draft-chunduri-lsr-isis-preferred-path-routing-07 (work in progress), November 2021.

15.2. Informative References

[I-D.bryant-ipfrr-tunnels]
Bryant, S., Filsfils, C., Previdi, S., and M. Shand, "IP Fast Reroute using tunnels", draft-bryant-ipfrr-tunnels-03 (work in progress), November 2007.

[I-D.ietf-rtgwg-segment-routing-ti-lfa]
Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", draft-ietf-rtgwg-segment-routing-ti-lfa-08 (work in progress), January 2022.

[I-D.xu-isis-encapsulation-cap]
Xu, X., Decraene, B., Raszuk, R., Chunduri, U., Contreras, L. M., and L. Jalil, "Advertising Tunnelling Capability in IS-IS", draft-xu-isis-encapsulation-cap-07 (work in progress), October 2016.

[RFC3936] Kompella, K. and J. Lang, "Procedures for Modifying the Resource reSerVation Protocol (RSVP)", BCP 96, RFC 3936, DOI 10.17487/RFC3936, October 2004, <https://www.rfc-editor.org/info/rfc3936>.

[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, <https://www.rfc-editor.org/info/rfc4090>.

[RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <https://www.rfc-editor.org/info/rfc5286>.

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <https://www.rfc-editor.org/info/rfc5714>.

[RFC5715] Shand, M. and S. Bryant, "A Framework for Loop-Free Convergence", RFC 5715, DOI 10.17487/RFC5715, January 2010, <https://www.rfc-editor.org/info/rfc5715>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC6981] Bryant, S., Previdi, S., and M. Shand, "A Framework for IP and MPLS Fast Reroute Using Not-Via Addresses", RFC 6981, DOI 10.17487/RFC6981, August 2013, <https://www.rfc-editor.org/info/rfc6981>.

[RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <https://www.rfc-editor.org/info/rfc7490>.

[RFC7812] Atlas, A., Bowers, C., and G. Enyedi, "An Architecture for IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)", RFC 7812, DOI 10.17487/RFC7812, June 2016, <https://www.rfc-editor.org/info/rfc7812>.

[RFC8333] Litkowski, S., Decraene, B., Filsfils, C., and P. Francois, "Micro-loop Prevention by Introducing a Local Convergence Delay", RFC 8333, DOI 10.17487/RFC8333, March 2018, <https://www.rfc-editor.org/info/rfc8333>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

[RFC8518] Sarkar, P., Ed., Chunduri, U., Ed., Hegde, S., Tantsura, J., and H. Gredler, "Selection of Loop-Free Alternates for Multi-Homed Prefixes", RFC 8518, DOI 10.17487/RFC8518, March 2019, <https://www.rfc-editor.org/info/rfc8518>.

[RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/info/rfc8754>.

Authors' Addresses

Stewart Bryant
University of Surrey ICS

Email: sb@stewartbryant.com

Uma Chunduri
Intel

Email: umac.ietf@gmail.com

Toerless Eckert
Futurewei Technologies Inc

Email: tte+ietf@cs.fau.de