Network Working Group
Internet-Draft
Intended status: Informational
Expires: 3 September 2026

F. Clad, Ed.
C. Filsfils
Cisco Systems, Inc.
Y. Su
D. Cai
Alibaba

2 March 2026

Efficient Remote Protection
draft-clad-rtgwg-efficient-remote-protection-00

Abstract

This document specifies Efficient Remote Protection (ERP), a mechanism for IP Fast Reroute (IP-FRR) that utilizes network notifications to activate pre-computed backup paths at nodes multiple hops upstream of a failure. ERP addresses scenarios where local protection mechanisms, such as Loop-Free Alternates (LFA) or Topology-Independent LFA (TI-LFA), result in suboptimal paths, specifically traffic hairpinning.

By activating protection at strategically selected upstream nodes rather than at the node immediately adjacent to the failure, ERP preserves routing optimality and prevents bandwidth waste. ERP applies to both complete link/node failures and performance degradations such as congestion or reduced link capacity. This makes ERP particularly beneficial in networks with high link utilization, such as AI data centers and Data Center Interconnect (DCI) networks.

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 3 September 2026.

Table of Contents

1. Introduction

IP Fast Reroute (IP-FRR) mechanisms ([RFC5714]) enable rapid traffic rerouting upon link or node failures by pre-computing and installing backup paths. Traditional IP-FRR mechanisms perform protection at the Point of Local Repair (PLR), which is the node directly adjacent to the failed resource. While this local protection model provides robust and immediate failure response, it has limitations:

* Topology Dependency: Loop-Free Alternates (LFA, [RFC5286]) and Remote LFAs (RLFA, [RFC7490]) do not provide complete protection coverage in all topologies, as described in [RFC6571].

* Path Optimality: While Topology-Independent Loop-Free Alternate (TI-LFA, [RFC9855]) provides complete protection coverage and enforces the post-convergence path from the PLR's perspective, it may steer traffic through suboptimal paths from the perspective of upstream nodes. Specifically, when the PLR's backup path traverses nodes that are upstream of the PLR on the primary path, traffic originating from or transiting through those nodes experiences hairpinning, where packets flow toward the PLR before being redirected back toward the destination.

Efficient Remote Protection (ERP) addresses the path optimality limitation by introducing a notification-triggered protection model. ERP allows strategically selected upstream nodes to pre-compute and install backup paths that protect against failures and performance degradations multiple hops away and activate these paths upon receiving a network notification ([I-D.ietf-rtgwg-net-notif-ps]) from the node adjacent to the affected resource. For complete failures, ERP reroutes all traffic; for performance degradations, ERP may load-balance traffic across primary and backup paths. This approach enables traffic to be rerouted from the most efficient location(s) in the network, avoiding hairpins and preserving optimal routing under failure and degradation conditions.

ERP is particularly beneficial in networks with high link utilization, such as those supporting AI workloads, where traffic patterns are highly synchronized, flows are large, and link capacity must be used efficiently.

2. Terminology

This document uses the following terms:

* Point of Local Repair (PLR): The node directly adjacent to a failed resource (link or node). In traditional IP-FRR mechanisms, the PLR is responsible for activating backup paths.

* Point of Remote Repair (PRR): A node that is one or more hops upstream of the PLR and that activates a pre-computed backup path upon receiving a network notification about a remote failure.

* Protected Resource: A link or node for which backup paths are computed and installed.

* Network Notification: A signal sent by a node upon detecting a local failure or performance degradation, intended to trigger protection mechanisms at remote nodes. Network notifications are described in [I-D.ietf-rtgwg-net-notif-ps].

* Hairpin: A suboptimal routing condition where traffic is forwarded toward a destination via a node that is located further from the destination than the traffic’s current location, resulting in unnecessary bandwidth consumption.

* Q-Space: The set of nodes from which a destination can be reached without traversing a given failed resource. This term is defined in [RFC9855].

3. Backup Path Efficiency

This section illustrates the path efficiency problem that ERP addresses through two scenarios.

Consider a generic scenario where a node R is adjacent to a resource (link or node) X, and traffic destined to D normally traverses X. Node R must protect destination D against the failure of resource X.

3.1. Local Protection with LFA

If node R has a directly attached LFA neighbor Q for destination D with respect to resource X, as shown in Figure 1 and Figure 2, R installs Q as a backup next-hop for destination D and activates it upon failure of resource X. In this case, the backup path is guaranteed not to create a hairpin. A fundamental property of LFA is that the LFA neighbor Q is not upstream of X (and therefore not upstream of R) on the shortest path to D.

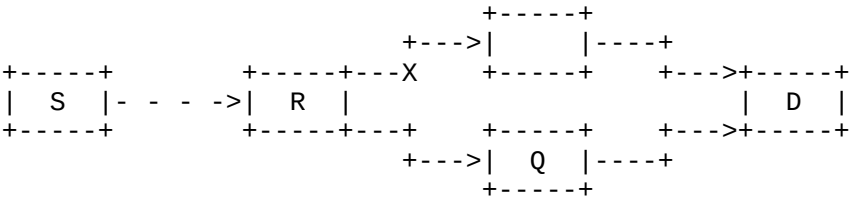

Figure 1: ECMP-LFA protection: R uses another ECMP path via Q

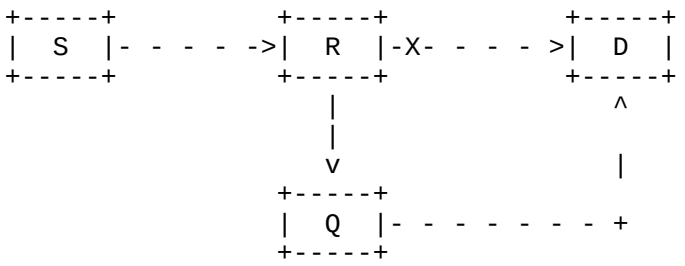

Figure 2: LFA protection: R uses directly attached neighbor Q

3.2. Local Protection with TI-LFA and Hairpinning

If node R does not have a directly attached LFA for destination D with respect to resource X, R may still be able to protect against the failure of X by using TI-LFA to enforce a path through one or more intermediate nodes U_0, U_1, ..., U_k before reaching a node Q in the Q-Space of D with respect to resource X. The Q-Space, defined in [RFC9855], is the set of nodes from which the path to D is unaffected by the failure of X.

However, these intermediate nodes (U_0, U_1, ..., U_k) that are outside the Q-Space are upstream of X, and may be upstream of R, on the primary path to D. When traffic originates at or transits through such a node U_i, it follows the primary path toward R. Upon reaching R, the traffic may be redirected via a TI-LFA backup path back through U_i and then onward through Q to reach D. This creates a hairpin, where traffic is unnecessarily transmitted from U_i to R and back, as depicted in Figure 3.

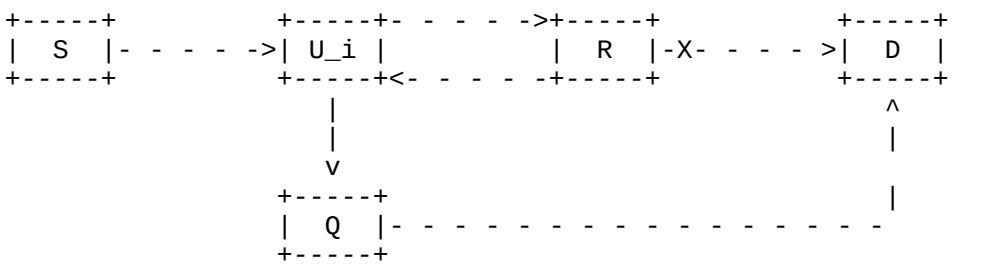

Figure 3: TI-LFA protection with hairpin: traffic from U_i to R and back

This hairpin consumes unnecessary bandwidth on the links between U_i and R in both directions, potentially causing congestion in networks with high link utilization, and increases end-to-end latency.

3.3. Remote Protection with ERP

The principle of ERP is to prevent hairpins such as the one depicted in Figure 3 by installing a backup path to D (protecting against the failure of resource X) at a node U that is upstream of R and would otherwise create a hairpin. When U receives a network notification from R indicating the failure of resource X, U activates its pre-installed backup path. This allows traffic originating at or transiting through U to be rerouted directly toward Q and then to D, without traversing R, as depicted in Figure 4.

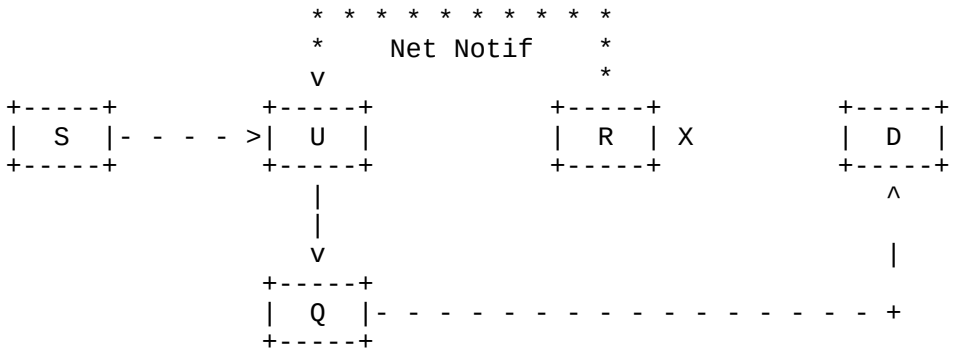

Figure 4: ERP protection: U activates backup path upon notification from R

The ERP backup path is enforced using Segment Routing ([RFC8402]) and encoded as a loop-free segment list from node U that steers traffic to a node in the Q-Space of D with respect to resource X, without traversing the failed resource X.

4. Backup Path Computation and Installation

This section describes the steps for computing and installing ERP backup paths. The specific algorithms for path computation and the protocol mechanisms for network notification subscription are outside the scope of this document.

4.1. Identifying Candidates for Remote Protection

For a given destination D and protected resource X adjacent to node R, a node U is a candidate Point of Remote Repair (PRR) if:

1. Node U is upstream of R on the primary path to D, and

2. The TI-LFA backup path computed by R for destination D (protecting resource X) would traverse U, creating a hairpin for traffic originating at or transiting through U.

ERP backup paths should be installed preferably at PRR candidates that are closest to the protected resource to limit the number of ERP backup paths required.

4.2. Network Notification Subscription

A PRR candidate node U subscribes to network notifications for resource X from node R. The mechanism for establishing this subscription depends on the specific network notification protocol used and is outside the scope of this document.

4.3. Backup Path Computation

The PRR node U computes a backup path to destination D that protects against the failure of resource X and does not create a hairpin. The backup path is encoded as a loop-free segment list that steers traffic to a node in the Q-Space of D with respect to resource X.

The segment list should terminate at the first node in the Q-Space along the backup path to minimize the length of the segment list and allow traffic to follow the regular shortest path from that point onward.

4.4. Backup Path Installation

The PRR node U installs the computed backup path in its forwarding plane and associates it with the network notification for resource X from node R.

5. Backup Path Activation

Upon receiving a network notification for resource X from node R, a PRR node U activates its pre-installed ERP backup path for destination D as follows.

5.1. Complete Failure

If the notification indicates a complete failure of resource X, node U immediately reroutes all traffic destined to D through the ERP backup path.

5.2. Performance Degradation

If the notification indicates a performance degradation of resource X (such as reduced link capacity, or congestion), node U may load-balance traffic between the primary path (via R) and the ERP backup path.

When load-balancing is employed, the load-balancing ratio should be determined based on the severity of the performance degradation indicated in the notification and may be adjusted dynamically as conditions change, based on subsequent notifications.

6. Operational Considerations

6.1. Incremental Deployment

ERP is designed to coexist with existing IP-FRR mechanisms such as LFA and TI-LFA. Traffic that passes through a PRR with an installed and activated ERP backup path will be protected at that upstream location, while traffic that does not pass through an ERP-enabled PRR will continue to be protected by traditional local protection mechanisms at the PLR.

This allows for incremental deployment of ERP in a network. Operators may initially deploy ERP at strategic nodes (such as those carrying high traffic volumes or those most susceptible to hairpinning) without disrupting existing protection schemes. Over time, ERP deployment can be expanded to additional nodes as needed.

6.2. Coordination with the PLR

The PLR (node R) must maintain its local protection mechanisms (e.g., TI-LFA backup paths) even when ERP is deployed at upstream nodes. This ensures protection for traffic that does not pass through any PRR, as well as providing a fallback mechanism.

6.3. Network Notification Reliability

The effectiveness of ERP depends on the reliable and timely delivery of network notifications from the PLR to PRR nodes. Operators should ensure that the network notification mechanism provides sufficient reliability and low latency to meet the protection requirements of the network.

If a PRR does not receive a notification within an expected timeframe after a failure (e.g., due to notification loss), traffic will continue to be protected by the PLR’s local protection mechanism, albeit potentially with hairpinning.

6.4. Capacity Planning

Operators should consider the capacity of backup paths when deploying ERP. While ERP avoids hairpinning and improves path efficiency, the backup paths themselves must have sufficient capacity to carry the redirected traffic without causing congestion.

7. Security Considerations

To be done.

8. IANA Considerations

This document does not require any IANA actions.

9. References

9.1. Normative References

[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/rfc/rfc8402>.

9.2. Informative References

[I-D.ietf-rtgwg-net-notif-ps] Dong, J., McBride, M., Clad, F., Zhang, Z. J., Zhu, Y., Xu, X., Zhuang, R., Pang, R., Lu, H., Liu, Y., Contreras, L. M., Mehmet, D., and R. Rahman, "Fast Network Notifications Problem Statement", Work in Progress, Internet-Draft, draft-ietf-rtgwg-net-notif-ps-00, 11 February 2026, <https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-net-notif-ps-00>.

[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/rfc/rfc5286>.

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

[RFC6571] Filsfils, C., Ed., Francois, P., Ed., Shand, M., Decraene, B., Uttaro, J., Leymann, N., and M. Horneffer, "Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks", RFC 6571, DOI 10.17487/RFC6571, June 2012, <https://www.rfc-editor.org/rfc/rfc6571>.

[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/rfc/rfc7490>.

[RFC9855] Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute Using Segment Routing", RFC 9855, DOI 10.17487/RFC9855, October 2025, <https://www.rfc-editor.org/rfc/rfc9855>.

Acknowledgements

Authors' Addresses

Francois Clad (editor)
Cisco Systems, Inc.
France
Email: fclad.ietf@gmail.com

Clarence Filsfils
Cisco Systems, Inc.
Belgium
Email: cf@cisco.com

Yuanchao Su
Alibaba
China
Email: yitai.syc@alibaba-inc.com

Dennis Cai
Alibaba
China
Email: d.cai@alibaba-inc.com