Understand LFA and Remote LFA IP Fast Reroute

Contents

Introduction

This document describes how the IP Fast Reroute (FRR) provides fast recovery methods in Label Distribution Protocol (LDP) based networks.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Background Information

This is a lot simpler to implement. Loop Free Alternate (LFA) is similar to Multiprotocol Label Switching (MPLS) FRR, for example, it pre-installs the backup next-hop into the forwarding plane. LFA’s do not introduce any protocol extensions and can be implemented on a per router basis, which makes it a very attractive option.

Understand MPLS

FRR Options:

Loop Free Alternate (LFA) FRR pre-computes a loop-free alternate path and installs into the forwarding place. LFA is calculated based on route in equality.

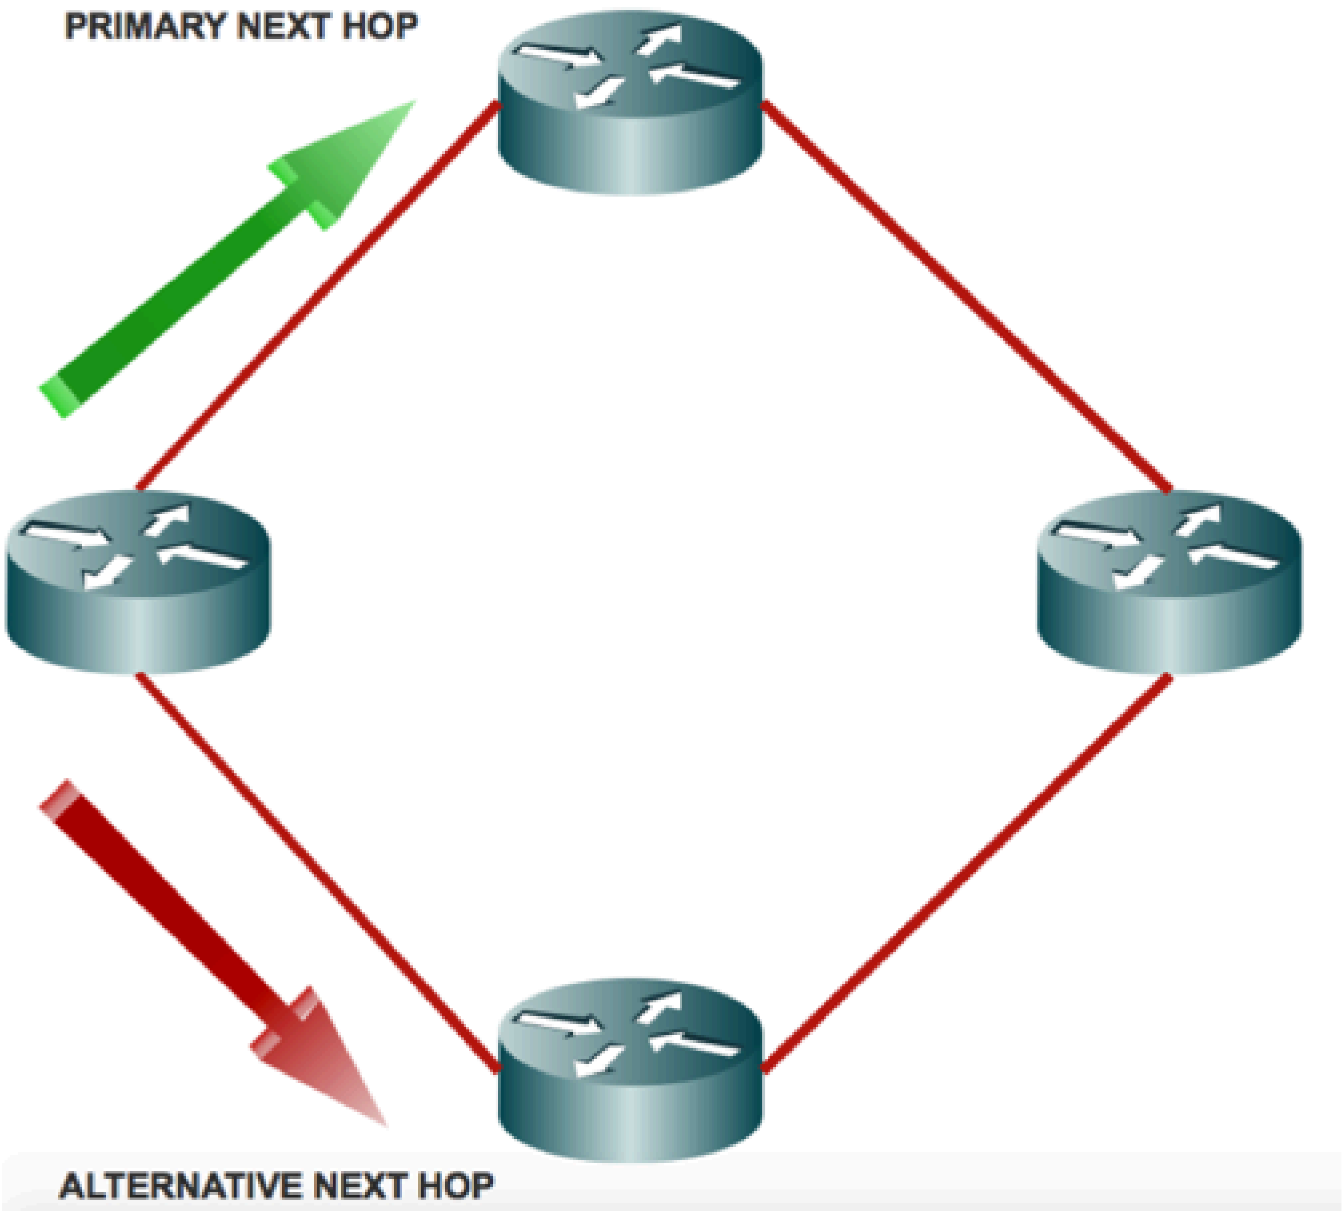

LFA:

Inequality 1: D(N,D) < D(N,S) + D(S,D)

Path is loop-free because N best path is not through local router. Traffic sent to backup next hop is not sent back to S.

Downstream Path:

Inequality 2: D(N,D) < D(S,D)

Neighbor router is closer to the destination than local router. Loop-free is guaranteed even with multiple failures (if all repair-paths are downstream path).

Node Protection:

Inequality 3: D(N,D) < D(N,E) + D(E,D) N path to D must not go through E.

The distance from the node N to the prefix via the primary next-hop is strictly greater than the optimum distance from the node N to the prefix.

Loop Free Link Protection for Broadcast Link:

Inequality 4: D(N,D) < D(N,PN) + D(PN,D)

The link from S to N must not be the same as the protected link.

The link from N to D must not be the same as the protected link.

Advantages of LFA and rLFA:

- Simplified Configuration
- Link and Node Protection
- Link and Path Protection
- LFA paths
- Support for both IP and LDP
- LFA FRR is supported with Equal Cost Multipath (ECMO)

Disadvantages of LFA and rLFA:

- LDP must be enabled everywhere
- Enabled Target LDP everywhere
- No other tunnel mechanisms other than MPLS are supported
- PQ Node protects only the link and not the node
- PQ Node calculations are only executed if there are unprotected paths for protectable Prefixes
- A targeted LDP session to PQ Node is only built if none exits yet
- No remote LFA for per-link

Remote LFA (rLFA):

LFA does not provide full coverage and it is very topology dependent. Reason is simple, for example, in many cases in order to backup next-hop, the best path goes through the router and calculates the backup next-hop.

This problem can be solved if you can find a router which is more than one hop away from the router that calculates, from which the traffic is forwarded to the destination that do not traverse the failed link and then you can tunnel the packet to that router.

These kinds of multi-hop repair paths are more complicated than single hop repair paths as computations are needed to determine if a path exits (to begin with) and then a mechanism to send the packet to that hop.

Look at a Point of Presence (POP) with aring topology as per the mentioned ring structure.

R3 does not meet inequality # 1 (3 < 1 + 2). So R3 best path is through the failed link.

If you find a node from which traffic is forwarded to the destination that do not traverse the failed link and it sends it to that node, then you can achieve FRR that does not cause a loop.

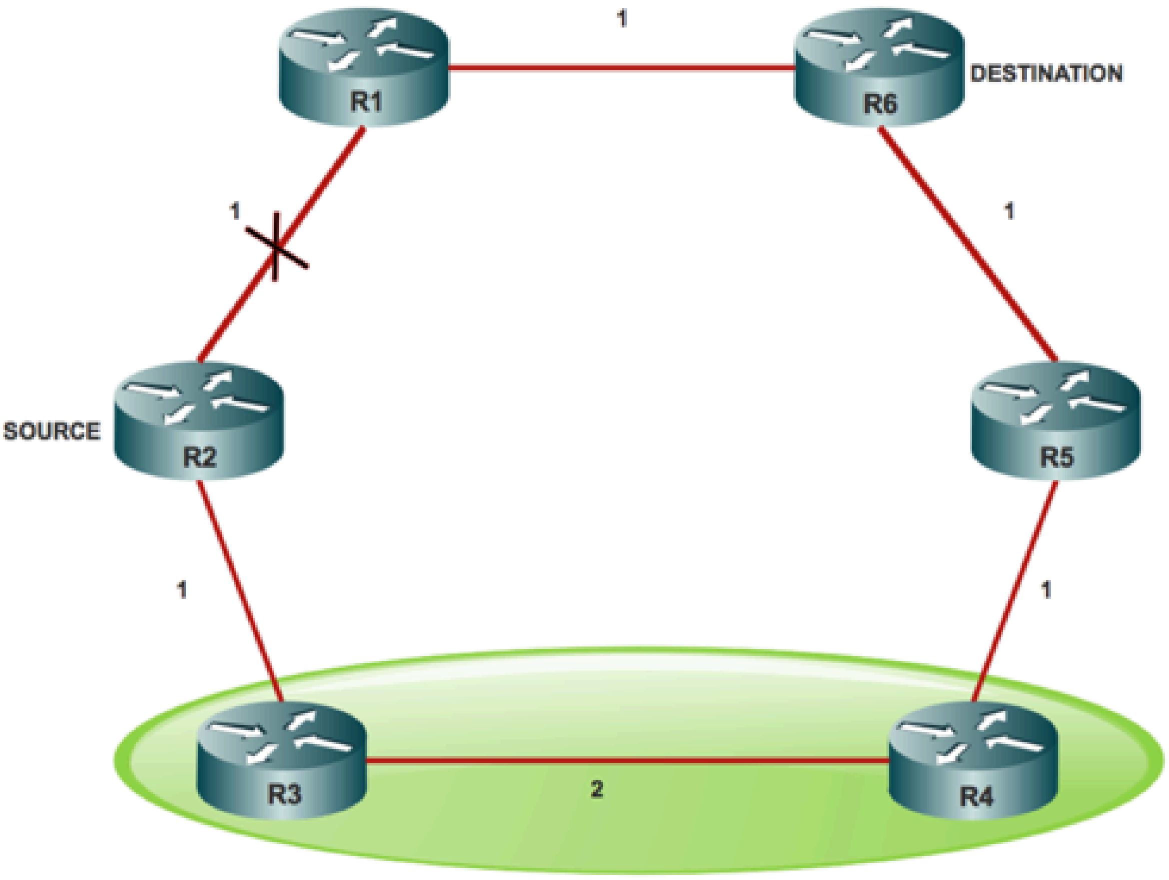

P-Space:

The P-Space of a router with respect to a protected link is the set of routers reachable from that specific router with the use of the pre-convergence shortest paths, without any of those paths, that transits that protected link.

P-Space is a set of routers that R2 (source) can reach without the use of the R2 (S) - R1 link which is R3 (P-Space) and R4 (P-Space) nodes.

Extended P-Space:

The extended P-Space of the router that protects with respect to the protected link is the union of the P-Space of the neighbors in that set of neighbor, with respect to the protected link, which makes it the union of the P-Spaces of the neighbours in that set of neighbors with respect to the protected link.

Extended P-Space contains the routers that are R2 - direct neighbor, R3 - can reach without the use of the R2 - R1 link which is R4 and R5 node. Point behind Extended P-Space is that it helps to increase the coverage.

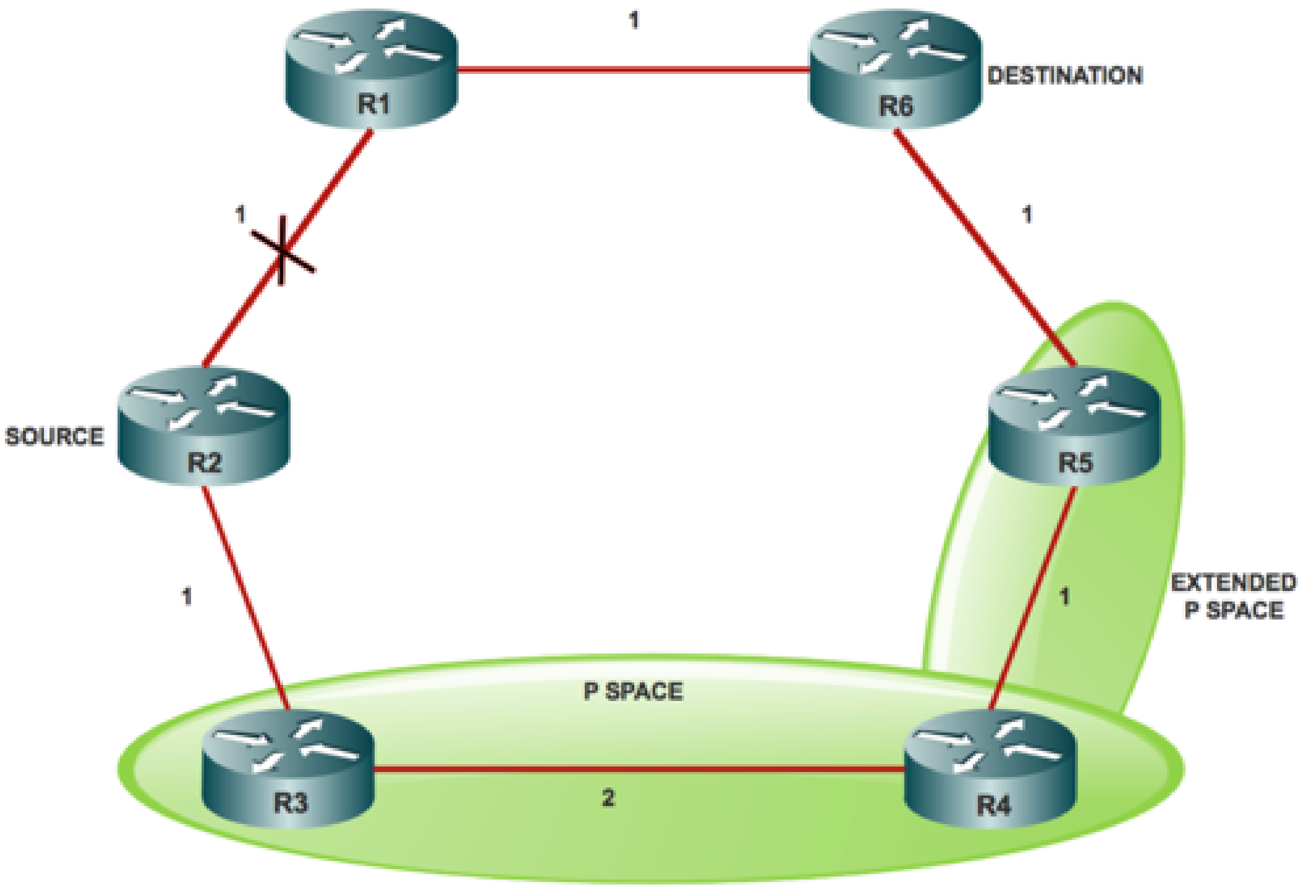

Q-Space:

Q-Space of a router with respect to a protected link is the set of routers from which that specific router that can be reached without any path (that includes ECMP splits) and transits that protected link.

Q-Space contains the routers that normally reach R6 without the use of the R2 (S) R1 link which is R1, R5 and R4 nodes.

PQ Node:

A router that is both Extended P-Space and Q-Space is a PQ node.

Any router which is a PQ node can be a remote LFA candidate. The candidate router to whom R2 (S) can send the packet, forwards the packet to the destination and does not traverse through R2(S) R1 link. In this case, R4 and R5 are the PQ nodes and are considered remote LFA candidates for R2 (S).

There are various ways to tunnel the traffic such as IPinIP, GRE and LDP. However, the most common form of implementation is LDP tunnel.

In Case of IP Traffic Protection:

If you protect IP traffic, then R2 (s) pushes an LDP label on top of IP packet to reach R4 (assume R2 (S) picket R4) as a Remote LFA node. When R3 receives the packet, it forwards the packet to R4 as a plain IP packet because of normal PHP behaviour. When R4 receives the packet destined to R6 (D), it forwards the packet upstream towards R5 node.

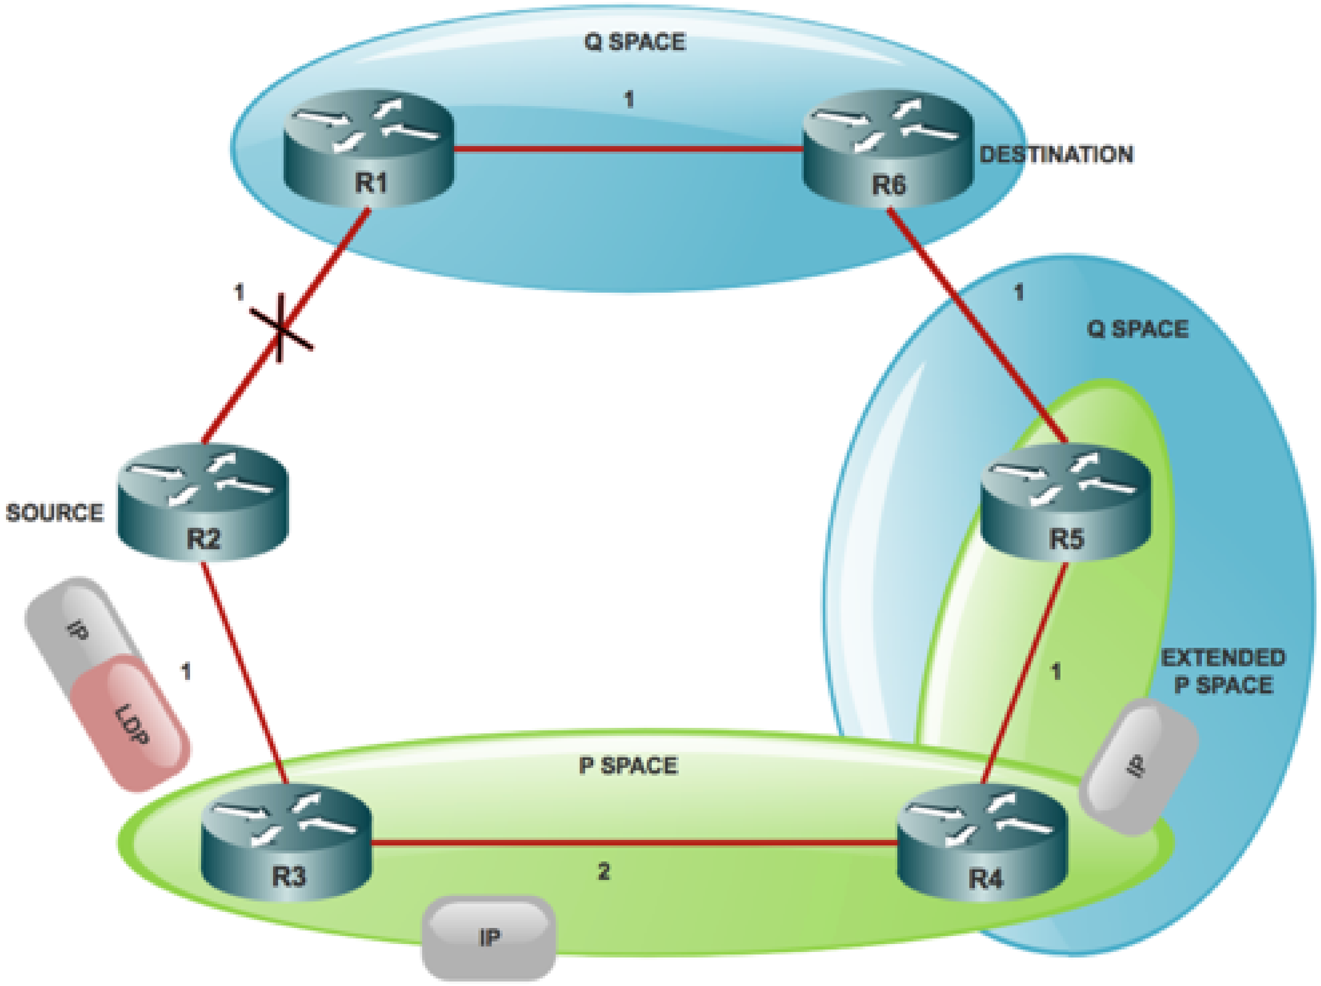

In Case of LDP Traffic Protection:

In this case a stack that consists of two LDP labels is used by R2(S).

Outer LDP label x, is the label to reach R4 and inner LDP label Y, is label to reach R6 (D) from R4.

Now the question is, how does R2 (S) know that R4 uses LDP label Y in order to send traffic towards R6(D). In order for the protective node-to-node to know what label a PQ node uses to forward the destination (D), it has to establish Targeted LDP session with aPQ node to get the FEC to label mapping. Therefore, you know that TLDP sessions must be enabled on the all the nodes for Remote LFA.

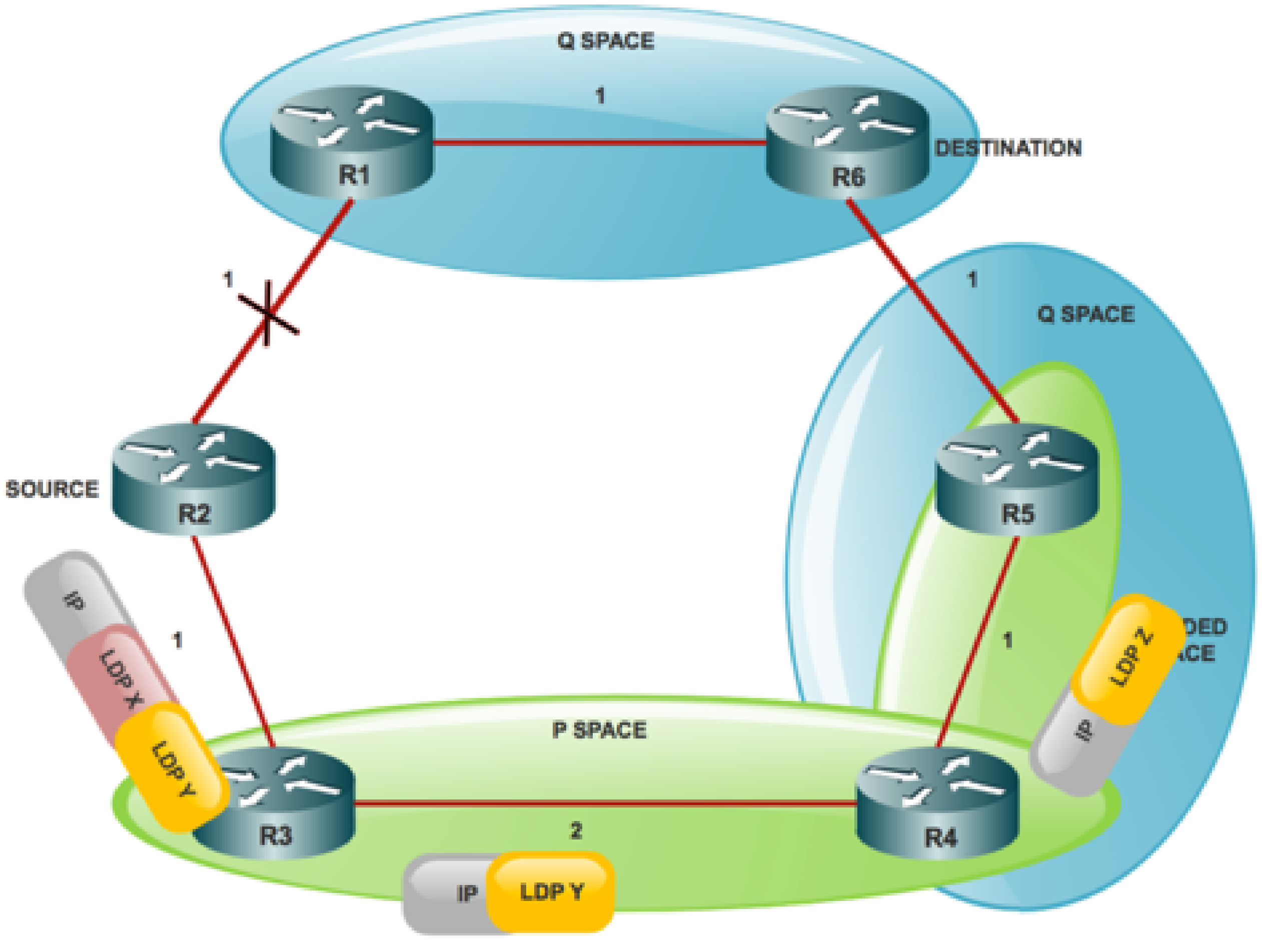

Benefits of rLFA over LFA:

- rLFA improves the LFA coverage in ring and poorly meshed topology
- It improves the consistency when the remote tunnel endpoint is selected
- Can work with RSVP with very little operational and computational overhead
- RSVP can be used to complement LFA/eLFA and vice versa
- When used in conjunction with MPLS LDP, there is no need of additional protocol in the control plane
- The data plane for MPLS makes use of label stack to tunnel the packets to the PQ node from there
- Traffic flows to the destination and does not return to the source or traverses the protected link

Configure

Network Diagram

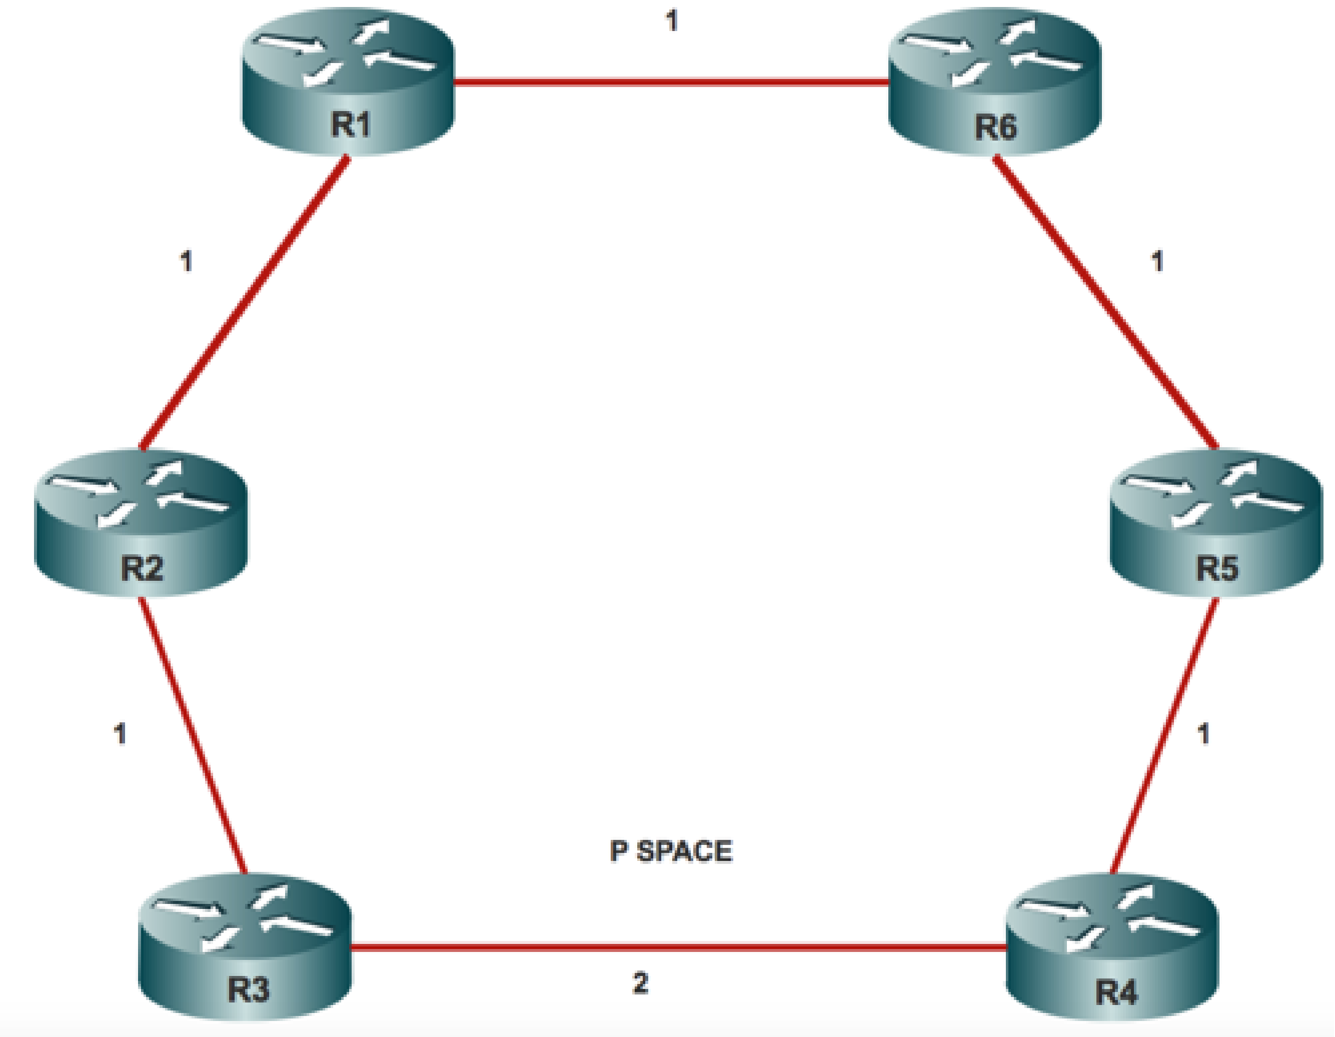

Configurations

Lab Details to Protect LDP Traffic:

ISIS Configuration:

```
router isis 20
net 20.0000.0000.0005.00
is-type level-1
metric-style wide level-1
fast-reroute per-prefix level-1 route-map LFA  >>>>>>>>>>> rLFA Configuration
fast-reroute remote-lfa level-1 mpls-ldp >>>>>>>>>>>>>>>>> rLFA Configuration
mpls ldp autoconfig level-1
```

MPLS Mandatory Configuration:

```
mpls ldp explicit-null
fast-reroute remote-lfa level-1 mpls-ldp
mpls ldp router-id Loopback0
```

Verify

Use this section in order to confirm that your configuration works properly.

In order to display the remote LFA tunnels for ISIS:

```
R1#show isis fast-reroute remote-lfa tunnels
Load for five secs: 0%/0%; one minute: 0%; five minutes: 0%
No time source, *11:28:59.528 UTC Wed Jan 3 2018
```

Tag 20 - FRR Remote-LFA Tunnels:

```
  MPLS-Remote-Lfa1: use Gi2/0, nexthop 10.3.4.4, end point 10.0.0.5
  MPLS-Remote-Lfa2: use Gi3/0, nexthop 10.3.3.3, end point 10.0.0.5
```

In order to check the Cisco IOS programming for a given prefix, run the CLI:

```
R1#show ip cef 10.0.0.5
Load for five secs: 0%/0%; one minute: 0%; five minutes: 0%
No time source, *11:32:04.857 UTC Wed Jan 3 2018

10.0.0.4/32
  nexthop 10.31.32.32 GigabitEthernet3/0 label [17|17]
    repair: attached-nexthop 10.3.4.4 GigabitEthernet2
  nexthop 10.3.4.4 GigabitEthernet2/0 label [17|17]
    repair: attached-nexthop 10.3.3.3 GigabitEthernet3
```

In this output, you can see primary and backup labels [17|17] respectively. The repair path goes via a remote LFA tunnel. It is not necessary that all the prefixes must be protected with the use of a remote LFA tunnel. Based on the possibility of looping, the LFA logic chooses to go over either normal backup path or a tunneled backup path.

```
R1#show ip route repair-paths 10.0.0.8
Load for five secs: 1%/0%; one minute: 0%; five minutes: 0%
No time source, *11:39:07.467 UTC Wed Jan 3 2018

Routing entry for 10.0.0.81/32

Known via "isis", distance 115, metric 30, type level-1
  Redistributing via isis 20
  Last update from 10.3.4.4 on GigabitEthernet2/0, 1d12h ago
  Routing Descriptor Blocks:
  * 10.3.4.4, from 10.10.0.81, 1d12h ago, via GigabitEthernet2/0
      Route metric is 30, traffic share count is 1
      Repair Path: 10.10.0.42, via MPLS-Remote-Lfa2
    [RPR]10.0.0.4, from 10.0.0.8, 1d12h ago, via MPLS-Remote-Lfa2
      Route metric is 20, traffic share count is 1
```

Troubleshoot

There is currently no specific troubleshoot information available for this configuration.