FAST REROUTE:
LFA, RLFA AND TI-LFA

Ron Bonica

NANOG 79

Juniper Networks
Engineering Simplicity

PROBLEM STATEMENT

- Link failures cause
  - Temporary loss of connectivity
  - IGP churn
- Connectivity is not be restored until the IGP converges
  - Milliseconds, seconds, or even tens of seconds
  - May impact user experience
- IP Fast Reroute
  - Precomputes repair paths that route around protected resources
  - Shifts traffic to the repair path immediately after a protected resource fails
  - Preserves connectivity during the IGP convergence period

IGP CONVERGENCE WITHOUT FAST REROUTE

- R2 calculates the least-cost route to R7
  - Via Link R2->R5
- Link R2->R5 breaks
  - R2 blackholes traffic destined for R7
  - Blackholing persists for milliseconds, seconds or tens or seconds until R2 calculates and installs a new least-cost route to R7
- Terminology
  - R2 is in a *pre-convergence* state regarding R7 from the time that Link R2->R5 breaks until it installs a new route to R7
  - R2 is in a *post-convergence* state regarding R7 after it installs a new route to R7

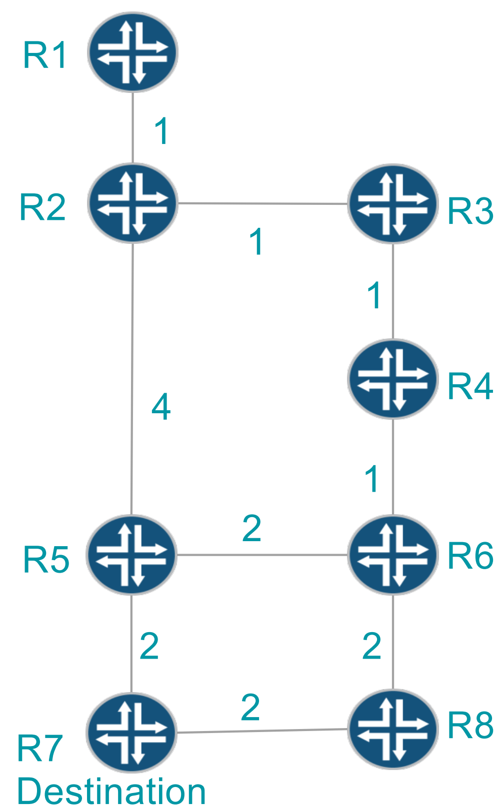

FAST REROUTE SOLUTIONS

- Loop Free Alternates (LFA)
  - Most severe restrictions on repair node selection
  - No tunnels required
  - Good coverage
- Remote LFA (RLFA)
  - Less severe restrictions on repair node selection
  - Tunnels required
  - Better coverage
- Topology Independent LFA (TI-LFA)
  - Least severe restrictions on repair node selection
  - Segment Routing (SR) tunnels required
  - Best (100%) coverage

Loop Free Alternates (LFA)

LFA SUPPORTS FAST REROUTE

- LFA preserves connectivity during convergence by pre-computing and installing a backup route to the destination
- When the primary route becomes unusable, traffic shifts to the backup route

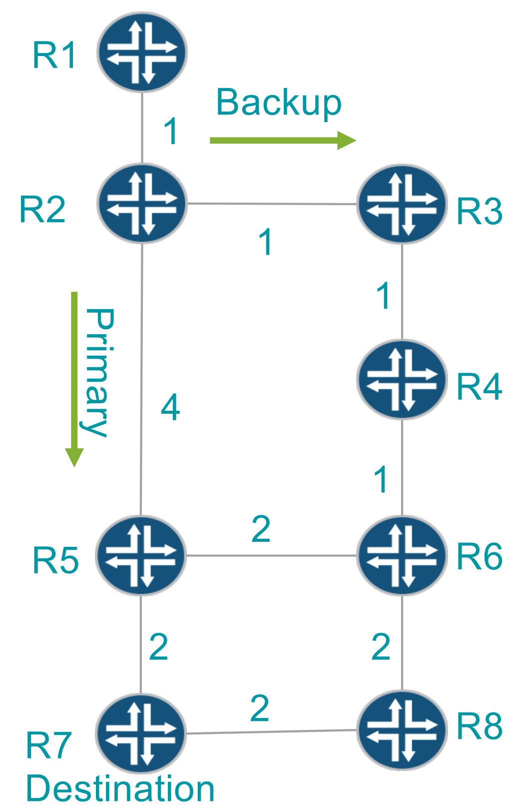

BEFORE THE BREAK

- LFA protects Link R2->R5
- R2 calculates the least-cost route to R7 via Link R2->R5
  - Installs as primary route to R7
- R2 identifies a repair node (i.e., R3)
- R2 identifies least-cost route to repair node via Link R2->R3
  - Installs as backup route to R7
- Terminology
  - Point of Local Repair (PLR) is the node upstream of the protected link
  - Repair node is the node to which packets are redirected when the primary route becomes unusable

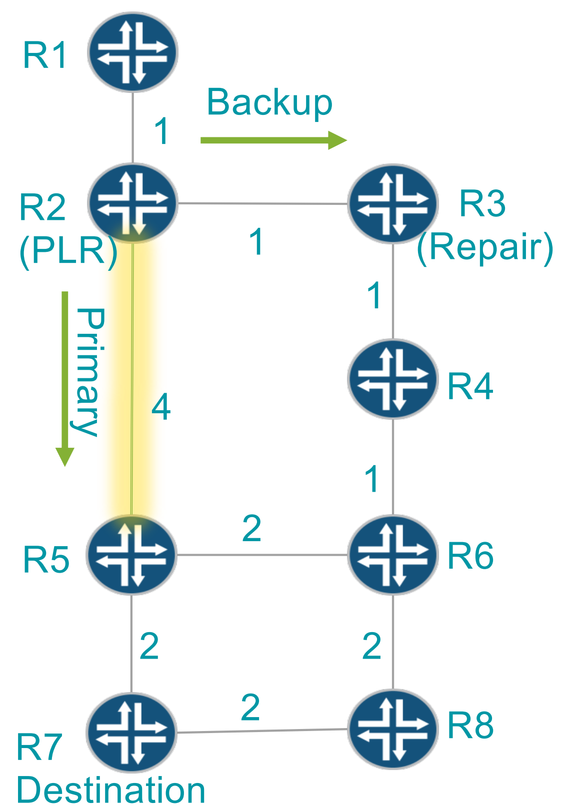

REPAIR NODE REQUIREMENTS

- Repair node must be directly connected to PLR
- Repair node must satisfy the *Loop Free Criterion*
  - When the repair node is in its pre-convergence state regarding the destination, its least-cost path to the destination must not traverse the PLR
- Repair node must satisfy the *Downstream Criterion*
  - When the repair node is in its pre-convergence state regarding the destination, its total cost to the destination must be less than the PLR's total cost to the destination

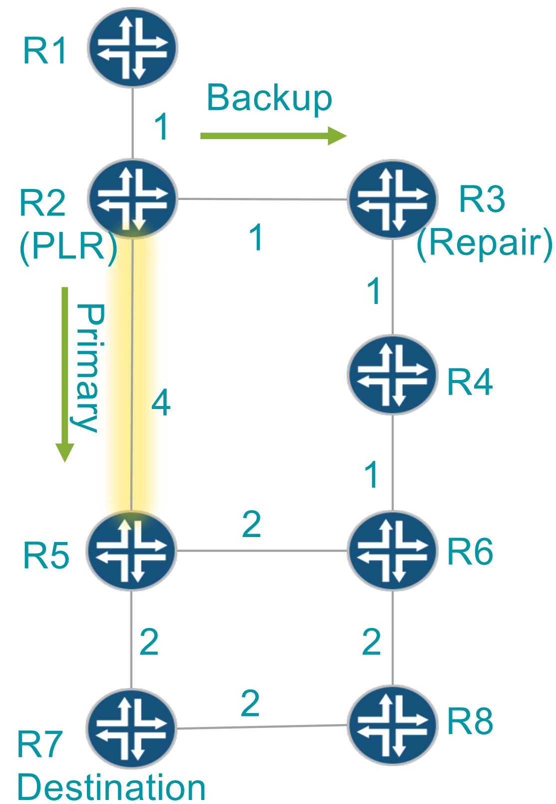

AFTER THE BREAK

- Link R2->R5 breaks
- R2's primary route to R7 (via Link R2->R5) becomes unusable
- R2 sends traffic destined for R7 via the backup route
  - That is, via Link R2->R3
- R3 sends traffic to R7 via R4
- Connectivity between R2 and R7 is preserved during convergence

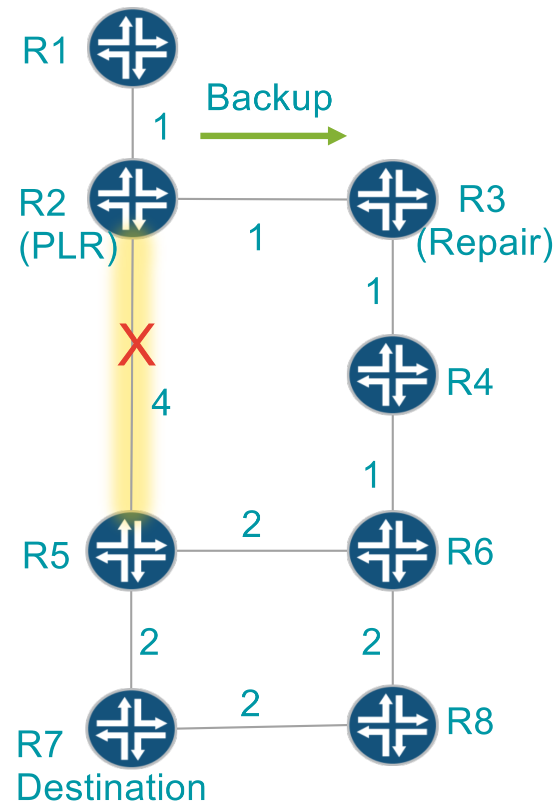

LFA LIMITATIONS

- LFA does not provide coverage for all destinations in all network topologies
- In the example:
  - LFA protects Link R2->R5
  - The IGP metric associated with Link R3->R4 is 3
- R3 does not satisfy the Loop Free Criterion regarding R7
  - Because in its pre-convergence state, its least-cost path to the destination traverses the PLR
- Therefore, in this topology, LFA cannot protect traffic flowing from R2 to R7 when Link R2->R5 fails

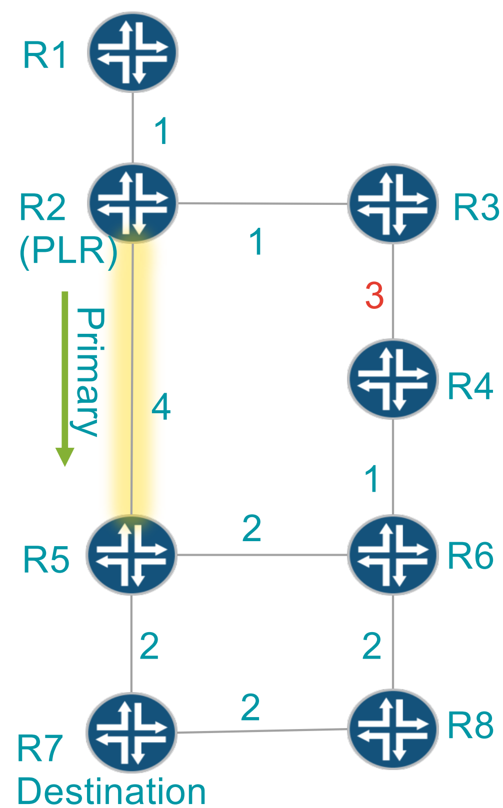

Remote Loop Free Alternates (RLFA)

RLFA ADDRESSES LFA LIMITATIONS

- RLFA addresses LFA limitations by allowing some nodes to serve as the repair node, even if they are not directly connected to the PLR
- A repair tunnel connects the PLR to the repair node
  - The repair tunnel follows the least-cost path from the PLR to the repair node
- RLFA supports link and node protection

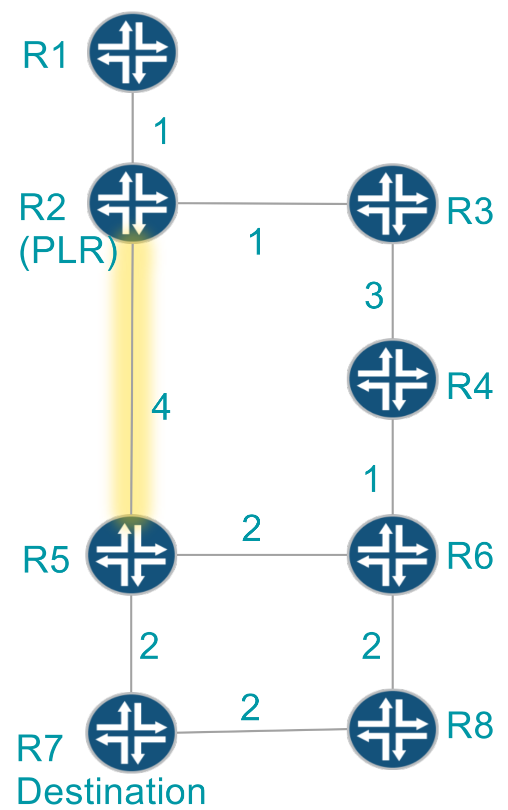

LINK PROTECTION: BEFORE THE BREAK

- RLFA protects Link R2->R5
- R2 calculates least-cost route to R7 via Link R2->R5
  - Installs as primary route to R7
- R2 identifies a repair node (i.e., R4)
  - In this example, the repair node is not directly connected to the PLR
- R2 creates a repair tunnel to the repair node
  - Repair tunnel follows the pre-convergence least-cost path from the PLR to the repair node
- R2 installs a backup route to R7 via the repair tunnel

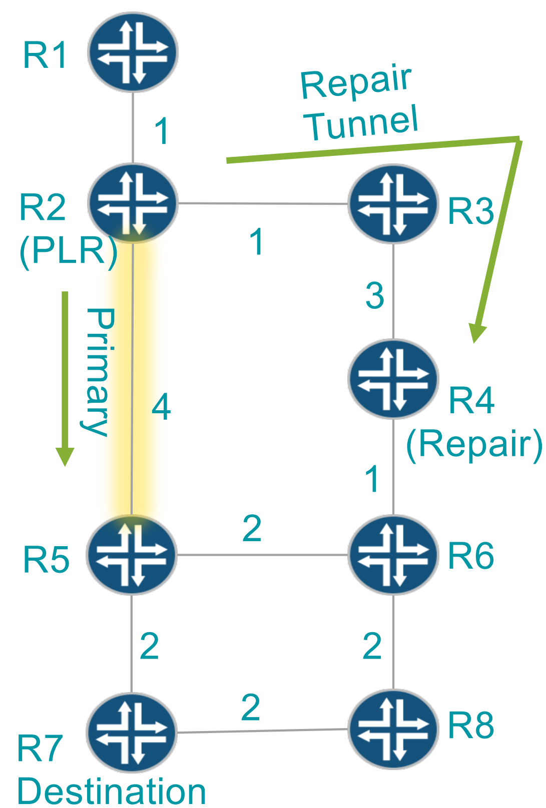

LINK PROTECTION: REPAIR NODE REQUIREMENTS

- The repair node must reside in PQ-space
  - That is, the repair node must reside in the intersection of the PLR's P-space and the destination's Q-space
- The PLR's P-space contains nodes that it can reach, in its pre-convergence state, *without traversing the protected link*
  - Nodes R3, R4, R6 and R8 are in P-space
- The destination's Q-space contains nodes that can reach the destination, in their pre-convergence state, *without traversing the protected link*
  - Nodes R4, R5, R6, R7 and R8 are in Q-space
- So, Nodes R4, R6 and R8 reside in PQ-space
  - R4 is chosen as the repair node, because it is closest to the PLR

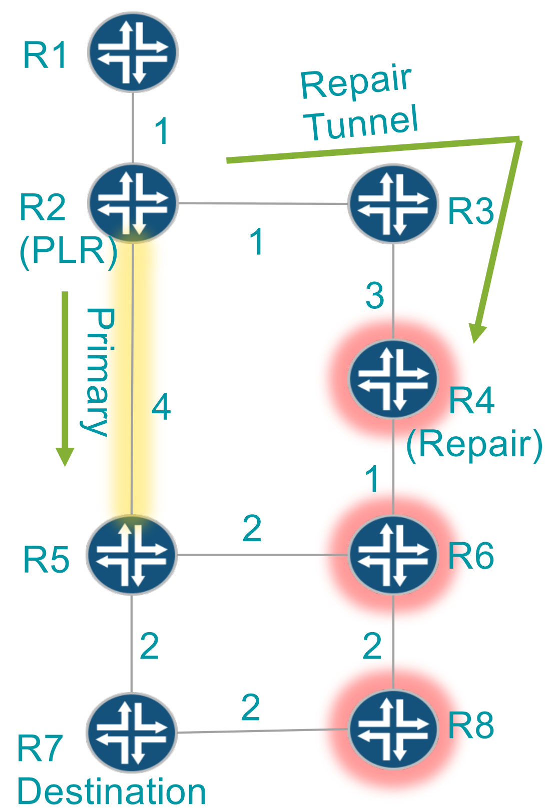

LINK PROTECTION: AFTER THE BREAK

- Link R2->R5 breaks
- R2's primary route to R7 (via Link R2->R5) becomes unusable
- R2 sends traffic destined for R7 via the backup route
  - That is, through the repair tunnel to R4
- R4 releases traffic from the tunnel and sends it to R7
  - Leveraging the ECMP between R6 and R7
- Connectivity between R2 and R7 is preserved during convergence

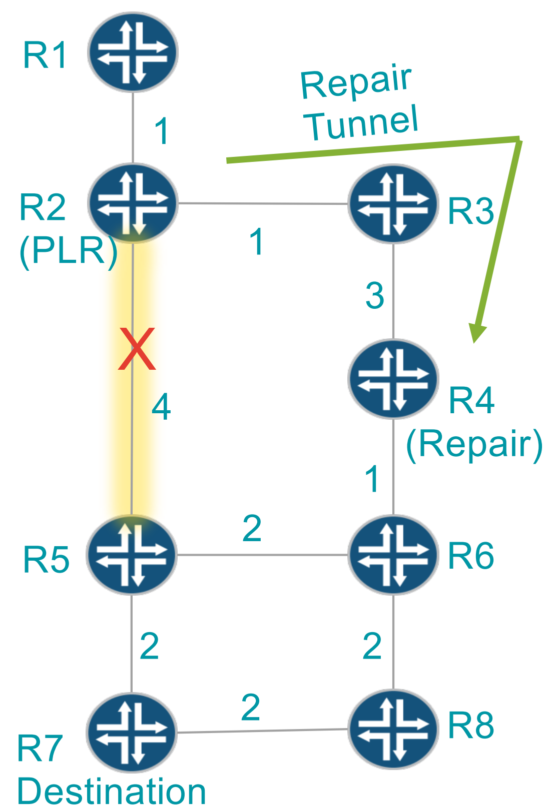

NODE PROTECTION: BEFORE THE BREAK

- RLFA protects Node R5
- R2 calculates least-cost route to R7 via Link R2->R5
  - Installs as primary route to R7
- R2 identifies a repair node (i.e., R8)
  - In this example, the repair node is not directly connected to the PLR
- R2 creates a repair tunnel to the repair node
  - Repair tunnel follows the pre-convergence least cost path to the repair node
- R2 installs a backup route to R7 via the repair tunnel

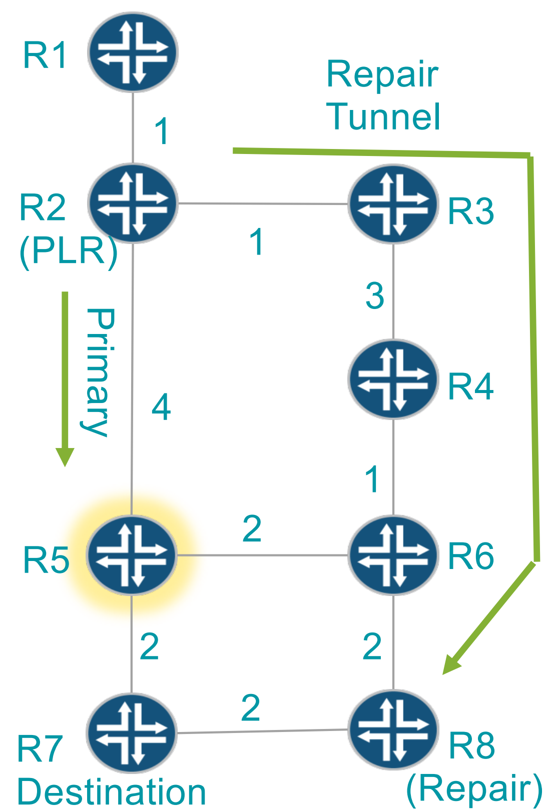

NODE PROTECTION: REPAIR NODE REQUIREMENTS

- The repair node must reside in PQ-space
  - That is, the repair node must reside in the intersection of the PLR's P-space and the destination's Q-space
- The PLR's P-space contains nodes that it can reach, in its pre-convergence state, *without traversing the protected node*
  - Nodes R3, R4, R6 and R8 are in P-space
- The destination's Q-space contains nodes that can reach the destination, in their pre-convergence state, *without traversing the protected node*
  - Nodes R7 and R8 are in Q-space
  - R6 is not in Q-space because of the ECMP route through R5
- So, only Node 8 resides in PQ-space

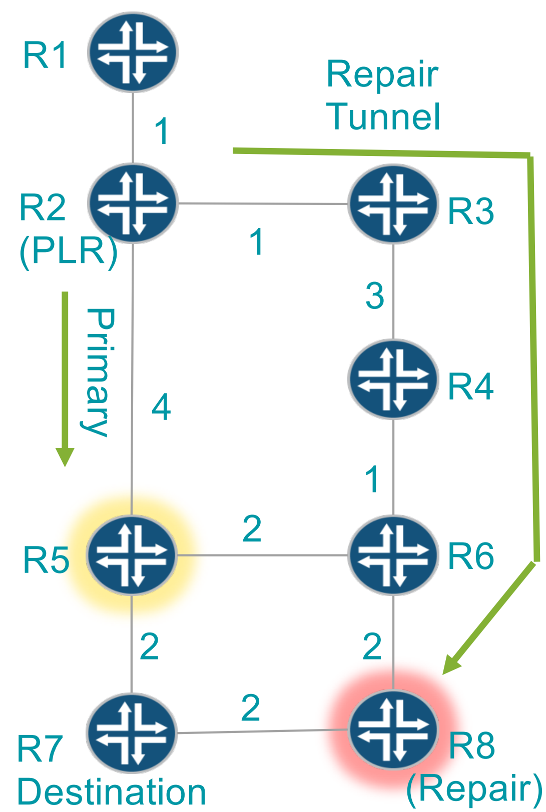

LFA LIMITATIONS

- RLFA does not provide coverage for all destinations in all network topologies
- In the example:
  - LFA protects Link R2->R5
  - The IGP metric associated with Link R4->R6 is 7
- R2's P-space contains only Node R3
- R7's Q-space contains Nodes R2, R4 and R5
- PQ-space is empty
- Therefore, in this topology, RLFA cannot protect traffic flowing from R2 to R7 when Link R2->R5 fails

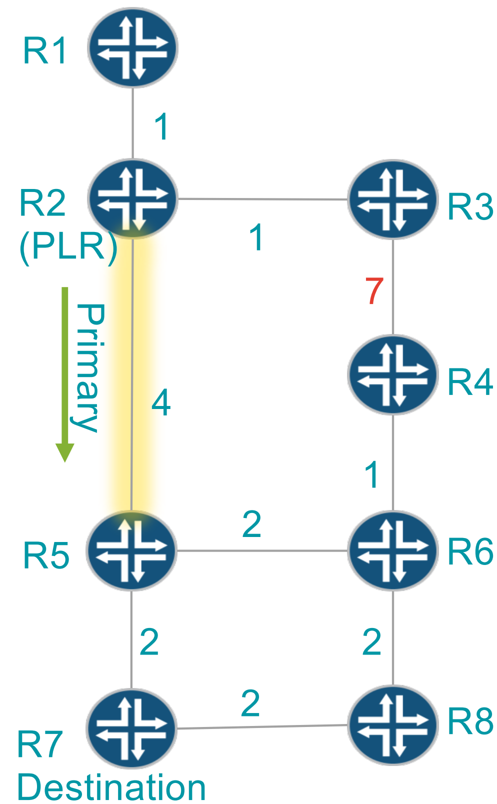

Topology Independent Loop Free Alternates (TI-LFA)

TI-LFA ADDRESSES RLFA LIMITATIONS

- TI-LFA addresses RLFA limitations by using SR paths as repair tunnels
  - Because the repair tunnel is an SR path, it is not required to traverse the least-cost path from the PLR to the repair node.
  - It can traverse any viable path
- So, the repair node can be outside of the PLR's P-space.
- However, the repair node must be within the destination node's Q-space.
  - Sometimes, the repair node is also the destination node
- Supports link and node protection

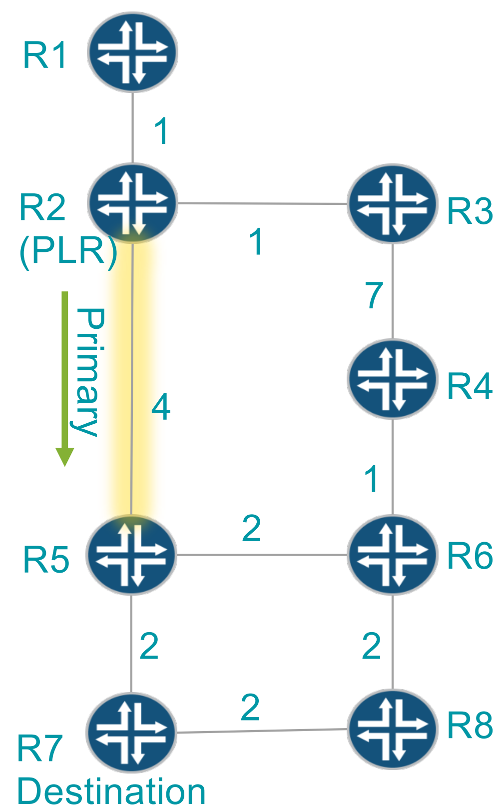

LINK PROTECTION: BEFORE THE BREAK

- TI-LFA protects Link R2->R5
- R2 calculates least-cost route to R7 via Link R2->R5
  - Installs as primary route to R7
- R2 identifies a repair node (i.e., R4)
- R2 creates an SR repair tunnel to the repair node
  - Repair tunnel has an outgoing PLR interface (i.e., Link R2->R3) and a repair list
  - The repair list includes an adjacency segment (MPLS label 1010) that begins on R3 and ends on R4
- R2 installs a backup route to R7 via the repair tunnel

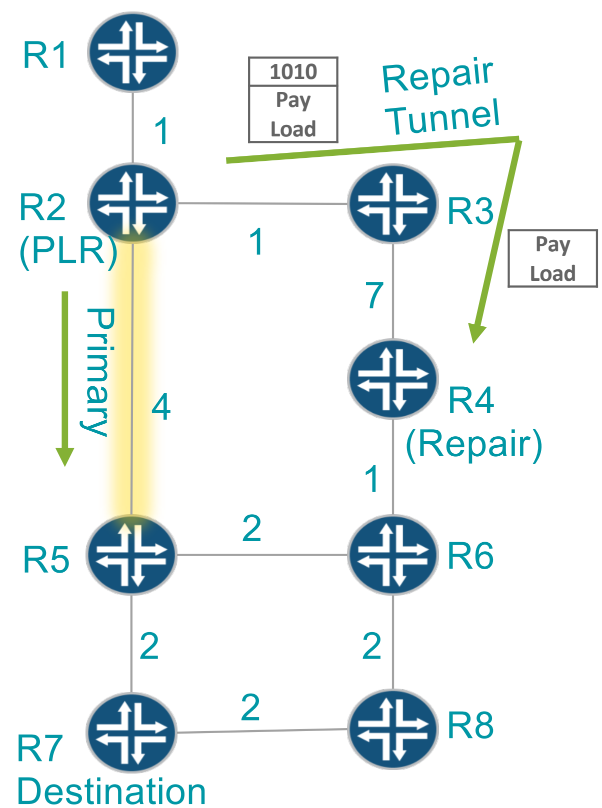

LINK PROTECTION: REPAIR NODE REQUIREMENTS

- The repair node must reside in the destination node's Q-space
- The destination node's Q-space contains nodes that can reach the destination, in their pre-convergence state, *without traversing the protected link*
  - Nodes R4, R5, R6, R7 and R8 are in Q-space
- R4 is chosen as the repair node because it is closest to the PLR

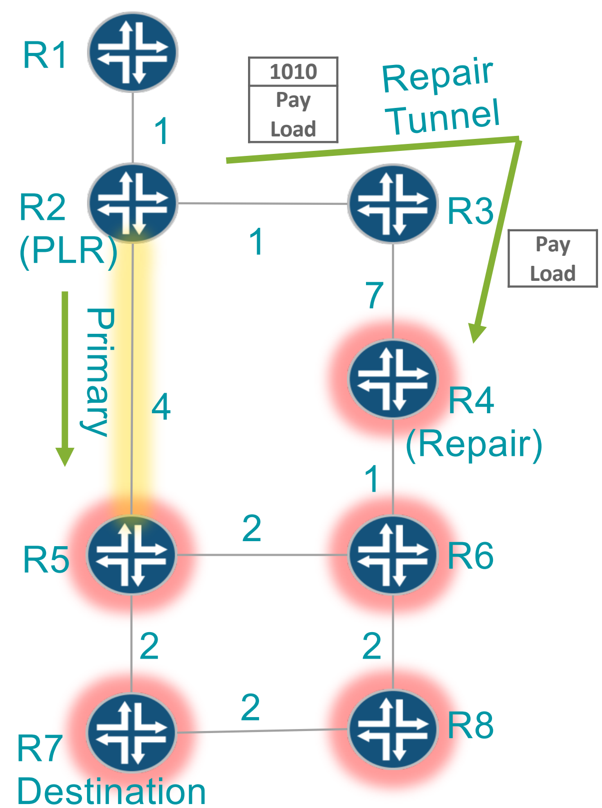

LINK PROTECTION: AFTER THE BREAK

- Link R2->R5 breaks
- R2's primary route to R7 (via Link R2->R5) becomes unusable
- R2 sends traffic destined for R7 via the backup route
  - R2 pushes MPLS label 1010 and forwards to R3
  - R3 pops MPLS label 1010 and forwards to R4
- R4 forwards packet along via least-cost route to R7
  - Leveraging the ECMP between R6 and R7
- Connectivity between R2 and R7 is preserved during convergence

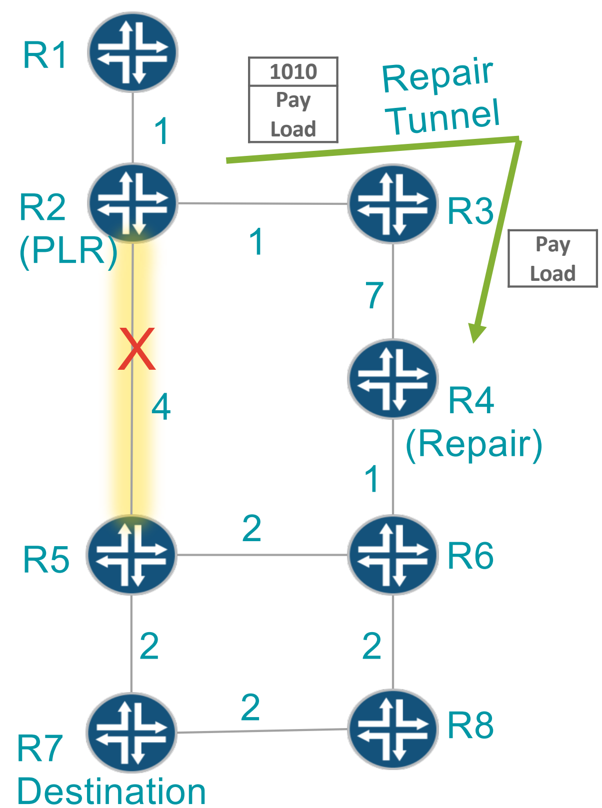

NODE PROTECTION: BEFORE THE BREAK

- TI-LFA protects Node R5
- R2 calculates least-cost route to R7 via Link R2->R5
  - Installs as primary route to R7
- R2 identifies a repair node (i.e., R8)
- R2 creates an SR repair tunnel to the repair node
  - Repair tunnel has an outgoing PLR interface (i.e., Link R2->R3) and a repair list
  - The repair list includes and adjacency segment (MPLS label 1010) that begins on R3 and ends on R4. It also includes a node segment (MPLS label 2010) that ends on R8.
- R2 installs a backup route to R7 via the repair tunnel

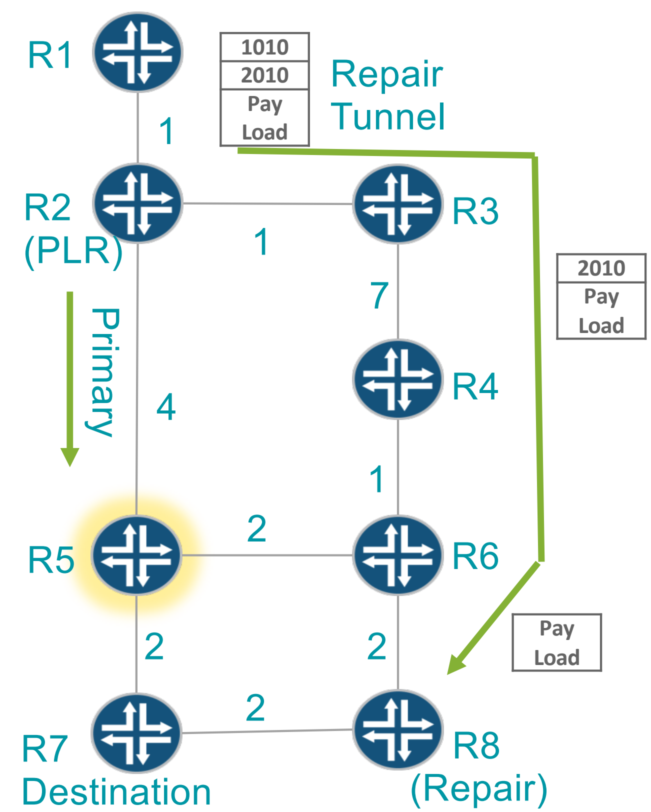

NODE PROTECTION: REPAIR NODE REQUIREMENTS

- The repair node must reside in the destination node's Q-space
- The destination node's Q-space contains nodes that can reach the destination, in their pre-convergence state, *without traversing the protected link*
  - R7 and R8 are in Q-space
  - R4 and R6 are not in Q-space because of the ECMP to R7 through R5
- R8 is chosen as the repair node because it is closest to the PLR

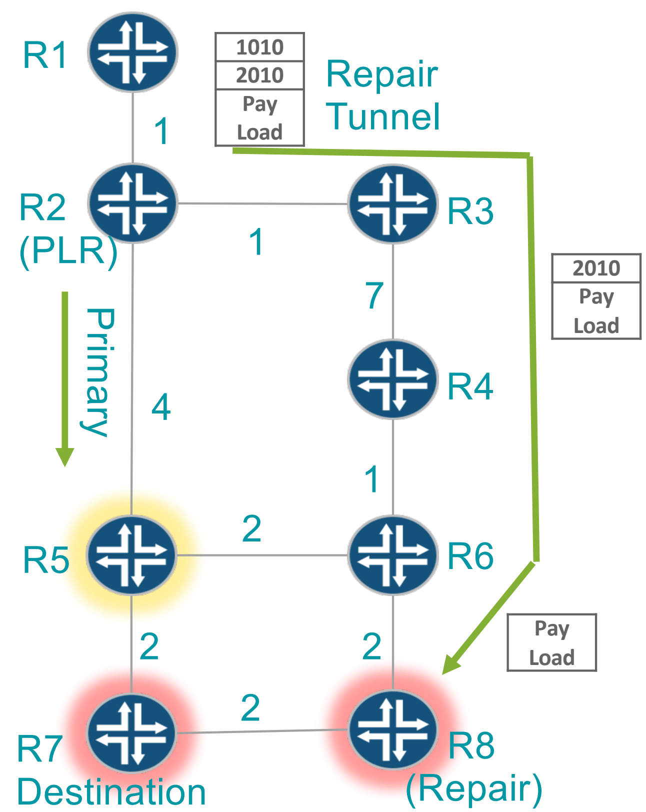

NODE PROTECTION: AFTER THE BREAK

- Link R2->R5 breaks
- R2's primary route to R7 (via Link R2->R5) becomes unusable
- R2 sends traffic destined for R7 via the backup route
  - R2 pushes MPLS labels 2010 and 1010 and forwards to R3
  - R3 pops MPLS label 1010 and forwards to R4
  - R6 pops MPLS label 2010 and forwards to R8
- R8 forwards packet along via least-cost route to R7
  - Avoiding R5
- Connectivity between R2 and R7 is preserved during convergence

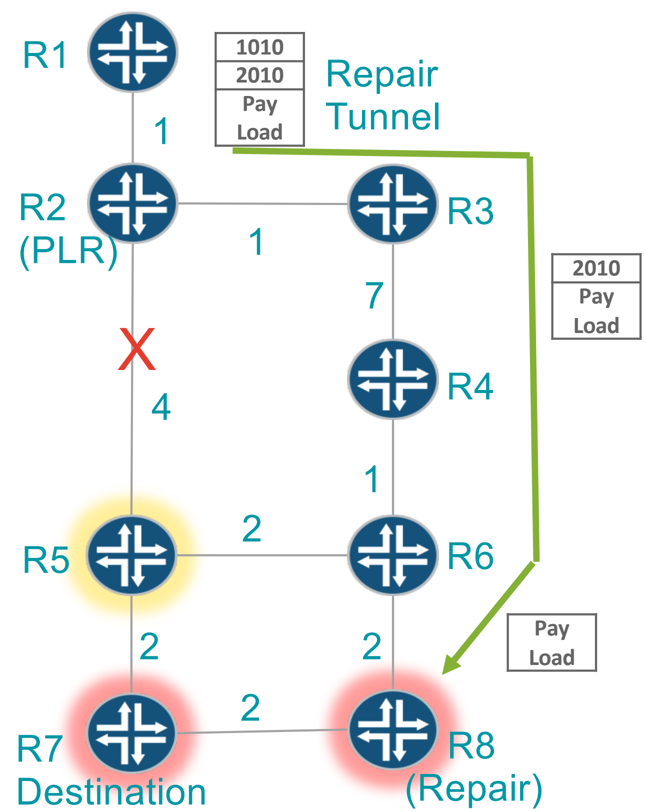

Conclusion

CONCLUSION

- TI-LFA is a simple FRR mechanism that can protect all nodes and all links in all network topologies
  - For this reason, we say that TI-LFA provides 100% coverage
- TI-LFA motivates many network operators to deploy SR
  - These operators want their traffic to traverse the least-cost path from source to destination
  - But require FRR in topologies that LFA and RLFA cannot support


THANK YOU
JUNIPER NETWORKS
Engineering
Simplicity