



**FOUNDRY
NETWORKS**

Adapting MPLS Fast Reroute to Create Resilient Rings

Ahmed Abdelhalim
Foundry Networks



MPLS JAPAN

www.mpls.jp

THE POWER OF PERFORMANCE



Ring Topology: Multiple Scenarios, Multiple Technologies



- ❑ Cost effective design
 - To be used when reducing fiber cost is paramount
 - Access or backbone applications
- ❑ Multiple technologies evolved over the years
 - FDDI
 - SONET Ring APS / SDH Ring MSP
 - Vendor proprietary ring resiliency for Ethernet rings
 - IEEE 802.17 RPR





Goals of a Good Ring Resiliency Mechanism

- ⊞ Highly survivable
 - Distributed control with no single point of failure
- ⊞ Sub-second resiliency
 - Ideally sub-50 ms resiliency
- ⊞ Provides traffic path optimization after link failure
- ⊞ Scalable
- ⊞ Can be easily applied/extended to multiple interconnected rings
- ⊞ Provides protection to upper layer services:
 - Routed services, VPN services, multicast-based services, etc.
- ⊞ Efficient bandwidth use (spatial reuse)
- ⊞ Simple to manage



Fast Reroute & Ring Topologies

What Makes Rings Different?



- ⊞ Fast Reroute mainly envisioned for rich meshed / partially meshed topologies
 - Potential for several backup paths
- ⊞ A ring topology does not offer rich connectivity. Only:
 - Primary path in one direction
 - Backup path has to be the other way around the ring
- ⊞ An FRR implementation with reasonable adaptation can efficiently offer resiliency for ring-based topologies



Why MPLS Fast Reroute for Rings?

- ⊗ Standards-based mechanism
- ⊗ Well established in the industry
 - Especially, for meshed topologies
- ⊗ Inherent rapid traffic protection
- ⊗ MPLS experiencing increasing popularity
 - Increasing demand for MPLS-capable hardware
- ⊗ Avoid running additional ring-specific resiliency mechanisms
 - Simplify whenever/wherever possible





Which Flavor of FRR?

- ⊗ Detour LSPs can easily allow for an implementation that offers
 - automatic protection
 - optimal forwarding paths under local repair and after path optimization
- ⊗ By extension, automatic bypass tunnels can also be used for higher scalability
- ⊗ Question: can FRR be used to efficiently provide ring protection?

YES



Optimal Calculation of Backup Paths Around the Ring





Legend

- iLER = ingress LER
- eLER = egress LER
- = Protected LSP
- - - = Backup Path (detour/bypass)

- ☛ CSPF should be used
- ☛ What may be a good heuristic for meshed topologies may not be good for rings
 - For a mesh topology, trying to find a backup path that flows towards downstream nodes of the protected path may work well
 - Yet for a ring, a backup path will have to flow in a **direction opposite** to the protected path
 - Excluding upstream nodes/links may make sense for loop avoidance in a mesh topology
 - Yet, upstream nodes/links **must not be excluded** for CSPF on a ring



Optimal Calculation of Backup Paths Around the Ring – cont.



- ☛ CSPF should optimize a **detour** (case B) to avoid awkward forwarding patterns (case A) in case of failure

- CSPF **primary goal** should be reaching eLER, not merely merging at NHop/NNHop
- CSPF should be recursive in finding the optimal merge point (MP): NHop, NNHop, ...
- eLER becomes the optimal merge point (MP) for the detour


- ☛ **Bypass tunnels** with automatic calculation can offer the scalability of many-to-one protection

- Algorithm must create an optimized bypass tunnel for each group of LSPs targeting the same eLER*!


* To avoid the undesirable case A!



Reducing Overhead and Increasing Scalability



- In theory, an LSP spanning N nodes can result in N-1 detours/bypass tunnels (case A)
- For detours, a ring offers many opportunities for straight forward detour merging
 - E.g., at each MP, merge detours having same SESSION, SENDER_TEMPLATE, outgoing interface, and next hop LSR*



- Merging can dramatically simplify the overall picture (case B)
 - Reduce signaling, state, and cross connects
- Reduce to practically **one** large detour spanning all nodes involved


Legend
MP = Merge Point



* An example of path-specific merging



Reducing Overhead and Increasing Scalability – cont.



All LSPs targeting eLER on this link are protected by one common bypass tunnel

- For bypass tunnels, many LSPs can be protected via the same bypass tunnel
 - On a link, protected LSPs targeting the same eLER can be mapped to one bypass tunnel





Ensuring Bandwidth Availability in Case of Failures

Straight Forward Approach

- ☛ Rely on spare capacity engineered into the ring
- ☛ Do not request bandwidth protection for the detours/bypass tunnels
 - It can mislead CSPF

A Variant of Above: A Differentiated Services Approach

- ☛ Business traffic has higher priority than residential traffic
- ☛ In case of failure, when capacity is reduced, business traffic is guaranteed* a percentage of the capacity at the expense of residential traffic



* Through traffic management / scheduling




11



Putting It All Together

What is the End Result Here?



- ☛ A plug-n-play resiliency mechanism
 - All backup paths set up automatically
- ☛ Rapid, sub-50 ms traffic protection via local repair → **ring wrap**
- ☛ Provides traffic path optimization after link failure
 - "Tunnel locally repaired" notification causes iLER to switch to the more optimal path → **ring steer**
- ☛ Provides protection to all MPLS-based services:
 - IP routing, pseudo-wire, VPLS, BGP/MPLS VPN, ...
- ☛ Combine with emerging MPLS OAM for improved visibility and manageability



12



Pros & Cons

Pros

- ☛ Reduce network complexity
 - One technology for enabling services, and offering resiliency
 - Avoid excessive and confusing resiliency at different layers
 - No lower layer technology operating in a ships in the night mode
 - No need to worry about interaction between those layers
 - Simplify training and operations
 - No need to learn another L-2 mechanism – like RPR, or RSTP, ...
- ☛ Operates on any MPLS capable hardware
 - No need for additional costly hardware provisions
- ☛ Reduce TCO
 - $TCO \propto k \cdot \text{Complexity}$
- ☛ More scalable and viable and than many would think

Cons

- ☛ Requires set up of the LSPs – not the backup paths
 - For large number of nodes (N), N^2 in a worst case*: each node shares service instances with every other node
 - An external tool, or mesh membership discovery can help rapidly set up those LSPs
- ☛ Each LSP requires maintaining state in the protected path direction, and its opposite direction



- * Classical MPLS problem. Not specific to the resiliency mechanism used
- * Using advertisement of membership in IGP (draft-ietf-ccamp-automesh)




13



Scoping the Viability of this Approach

The Metro Network




- ☛ Attractive solution for Metro Ethernet Networks
 - Can be easily extended to span multiple rings
- ☛ Contrary to common perceptions, an optimized implementation can offer
 - Well groomed signaling, paths, cross connects
 - Optimized paths and backup paths
 - Backup paths contained within their own rings → simplifies manageability



14



Scoping the Viability of this Approach *Beyond the Metro Network*



- ▣ It is recommended to break up the LSP and resiliency domains
 - Better manageability
 - Allow for autonomous control within the different markets/regions
- ▣ Choice of backbone resiliency mechanism is independent of the metro
 - May use same approach in case of rings
- ▣ Utilize emerging services that inherently support inter-domain connectivity
 - MS-PW, BGP/MPLS VPN, inter-AS VPLS



Summary

- ▣ MPLS Fast Reroute is a powerful and viable approach for providing ring resiliency
 - Provides rapid protection and optimization equivalent to the well known “Wrap then Steer” behavior
 - Can be set up with minimum effort – using today’s automatic detour/bypass tunnels
 - Far scalable than many would think
- ▣ Simplify overall design and operation
 - No need for an additional resiliency layer for handling ring resiliency
 - No need for an additional OAM layer – rely on MPLS OAM
- ▣ Reduce cost
 - Many service providers invest in MPLS capable hardware to offer next generation services
 - No need for special hardware designed specifically for addressing ring resiliency



**FOUNDRY
NETWORKS**

Thank You!

THE POWER OF PERFORMANCE

