

IP/LDP Fast-Reroute Using Maximally Redundant Trees draft-atlas-rtgwg-mrt-frr-architecture-01

Alia Atlas, Maciek Konstantynowicz, Robert Kebler,
Gábor Enyedi, András Császár,
Russ White, Mike Shand

IETF 82, Taipei, Taiwan

Goals Met

- **Guaranteed 100% Coverage** in an IGP area/level for **single link/node failure**.
- Provide fast-reroute for IP Unicast, LDP Unicast, IP Multicast and LDP Multicast
- Provide a live-live Multicast solution
- Support Incremental Deployment.

Solve the IP/LDP Fast-Reroute Problem Fully

Some Benefits of 100% Coverage

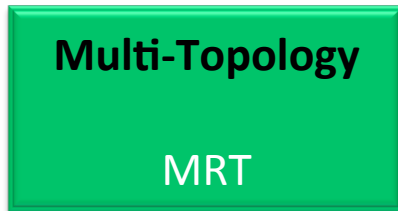
- Networks frequently have maintenance events and other changes – so network topology is frequently not the complete designed architecture.
 - Coverage is not dependent on specific architecture.
- Allows for Micro-Forwarding Loop Prevention techniques (which inevitably delay convergence) to be used without traffic loss.

Costs to Meet Goal

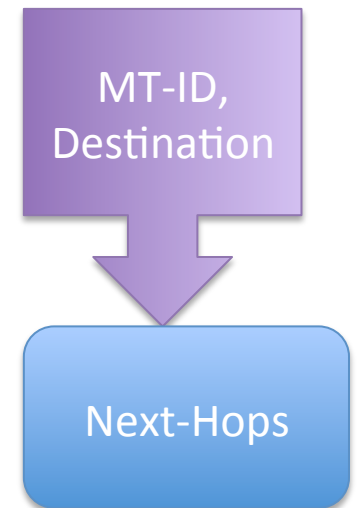
- Tag Packets to indicate *topology* being used
 - In MPLS networks, can basically allocate 2 additional labels per router loopback in area/level
- IGP Capabilities signaling
 - Learn what routers support MRT FRR for which forwarding and type of traffic.
- New Algorithm – MRTs computation
 - Computational complexity similar to a few SPFs
 - Routers need to implement same algorithm/result

Forwarding

To go beyond LFA, it is necessary to use an additional forwarding mechanism.



- Multi-Topology forwarding includes the topology (MT-ID) with the destination to find the next-hops.
- For each topology for each destination, need to compute next-hops.



MRT FRR: Multiple Forwarding Topologies – 1 Network Topology

Fast-Reroute with MRT uses 3 forwarding topologies:

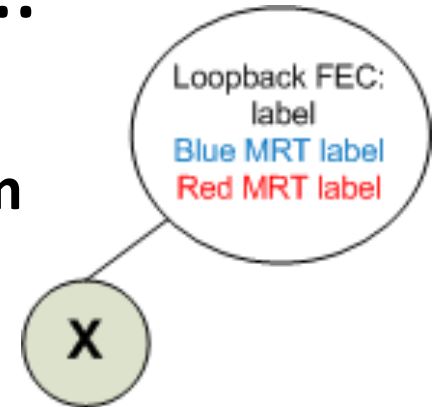
1. current default topology – next-hops computed by SPF
2. Blue MRT topology - MRTs computes next-hops
3. Red MRT - MRTs computes next-hops

Same network topology used as input to the SPF and MRT algorithms. Just different next-hops are computed as a result of different algorithms.

MRT Unicast Forwarding: For Want of 2 bits...

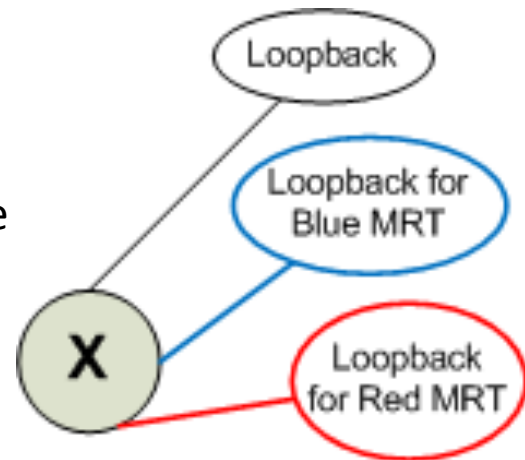
- **IP and LDP Unicast – Recommended Minimum**

- Simple and basic solution – no label-stacking needed
- MPLS label distributed by LDP to indicate FEC and MT-ID
- “Just works” with MPLS hardware.



- **Other Options Exist**

- Topology-ID label
 - Can reuse MPLS context-label space hardware
- IP-in-IP with different loopbacks



Use of MRT FRR for Unicast Fast-Reroute

- 1) Before failure:
 - a) Compute & Use Shortest-Path Tree Next-Hops
 - b) Compute MRTs and LFAs. Select alternates for each primary next-hop for each destination.
 - c) Install alternates into forwarding plane.
 - d) Install MRT next-hops into forwarding plane.
- 2) At failure: PLR moves traffic to alternates.
- 3) After failure:
 - a) Compute SPT next-hops and alternates. Install in forwarding plane.
 - b) Wait until network converged.
 - c) Install MRT next-hops into forwarding plane.

Incremental Deployment

- Before computing MRTs, prune out routers that don't support MRT FRR – determine the local island.
- Benefits possible as soon as two neighboring routers both support MRT FRR.
 - Easy way to make long rings work.

Next Application:

Multicast Fast-Reroute

- Traffic-Handling: A router only forwards alternate traffic when its upstream primary link(s) are down.
- Link-Protection: PLR replication into tunnels using unicast alternate to reach next-hops.
- Node-Protection: Do we need it?
 - a) PLR replication with tunnels
 - b) Distributed replication via alternate trees per (PLR, failure-point) – complexity and scaling...

Next Application: Multicast Live-Live

- Identify the MT-ID (Blue MRT or Red MRT) when signaling an (S,G) or mLDP.
 - Separate G for each MT-ID
 - Traffic can be distinguished on common link because of different G or MPLS label.
- Receivers join both the (S,G-blue) on the Blue MRT and (S,G-red) on the Red MRT.
 - Receivers determine which packets to keep
- MRT natural fit for Multicast Live-Live
 - MRTs rooted at the Multicast Source
 - Automatically computes maximally disjoint trees.

Summary

- Meets 100% coverage goals.
- Describes applications beyond unicast fast-reroute.
- Control-plane computation and alternate selection has been implemented for examining behavior on different topologies.
- Forwarding Plane is simple and no changes.
- Control Plane – complexity confined to deterministic algorithm.
- Good way forward to finish IP/LDP fast-reroute.

Ready to become a WG draft?