IP/LDP Fast-Reroute Using Maximally Redundant Trees

draft-ietf-rtgwg-mrt-frr-architecture-01

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

IETF 83, Paris, France

Status

- WG draft after Taipei
- Moved Multicast-related section to separate draft-atlas-rtgwg-mrt-mc-arch
- Focused on required forwarding mechanisms
- Added detail on IGP extensions, inter-area/ inter-level behavior, phased deployment, and multi-homed prefixes

Forwarding Mechanisms

- For LDP, MUST support single LDP label representing FEC plus MT-ID
  - Uses extensions from draft-ietf-mpls-ldp-multi-topology
- For IP, SHOULD support single LDP label representing FEC plus MT-ID
  - IP-in-IP also feasible

Multi-homed Prefixes for IP

- A proxy-node represents the multi-homed prefix and is attached to 2 routers in the graph.
- The Red MRT to proxy will go through the router X >> Proxy

  - So for IP traffic, the LDP label for (MT-ID=Red, X) or X's Red loopback address can be used for traffic destined to the Proxy and forwarded on the Red MRT.
- The Blue MRT to proxy will go through the router Y << Proxy
  - So the LDP label for (MT-ID=Blue,Y) or Y's Blue loopback address can be used.

Desired Inter-area/inter-level behavior

- For unicast fast-reroute, want to leave the blue/red MRT forwarding topology at the ABR/LBR and return to SPT
  - SPT uses shortest paths
  - Failures that appear in multiple areas (e.g. ABR/LBR failures) can be separately identified and repaired around
  - Packets can be fast-rerouted again in the next area/level, if necessary, due to a failure in a different area/level.
- How to get this behavior without additional/complex forwarding?

Inter-area/Inter-level Special Behavior

- IP-in-IP: traffic is sent to a loopback address meaning either (X, Red) or (Y, Blue)
  - When an ABR/LBR gets a packet destined to its Blue or Red loopback address, the outer IP header is removed and the packet is forwarded on the SPT. Trivially gives desired behavior.
- In LDP: traffic sent to (FEC, Red) or (FEC, Blue)
  - If FEC represents the ABR/LBR, then action is p affic is forwarded based on internal label/IP address
  - If FEC is for router A (in other area), then
    - ABR/LBR uses Label_A_SPT for the FEC in the ABR/LBR's best area/level
    - ABR/LBR provides Label_A_SPT for (FEC, Red) and (FEC,Blue) into the non-best area/level
    - ABR/LBR installs its best SPT next-hops for Label_A_SPT
    - Effect is equivalent to removing the MT marking on the way to the ABR/LBR

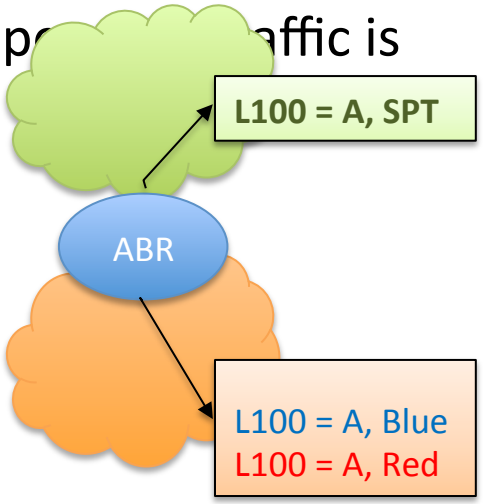

Good behavior – little extra computation – ABR/LBR decides and causes correct forwarding

MRT Phased deployment: 2 adjacent routers provide benefit

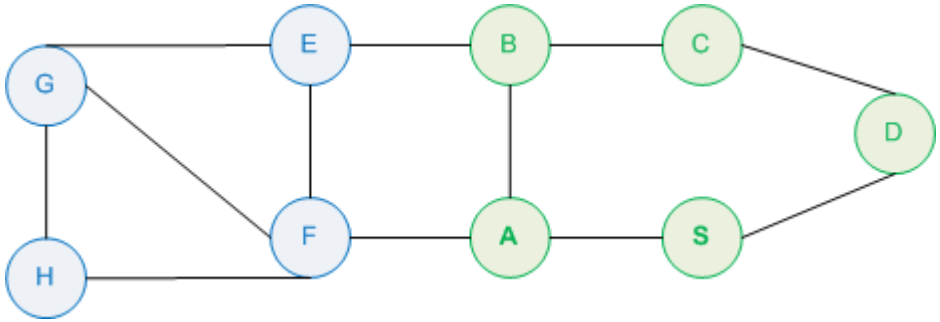

For a destination G, traffic from the sub-area can only go to adjacent loop-free nodes – so pick 2 closest to the destination G.

Then, create a proxy-node connected to the sub-area by the links of the 2 selected loop-free nodes.

For each destination, determine the adjacent loop-free nodes outside the MRT-supported sub-area to which traffic could be sent.

E and F are loop-free for destinations E,F, G and H with regard to the sub-area. SPT to G is shown.

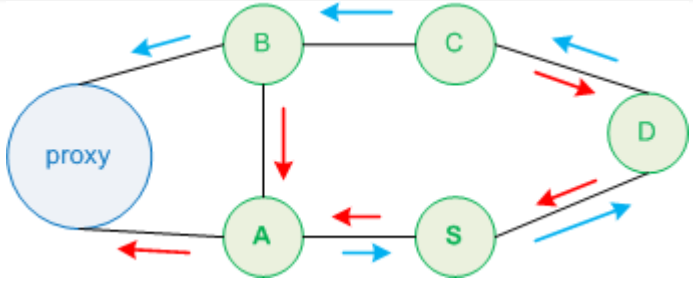

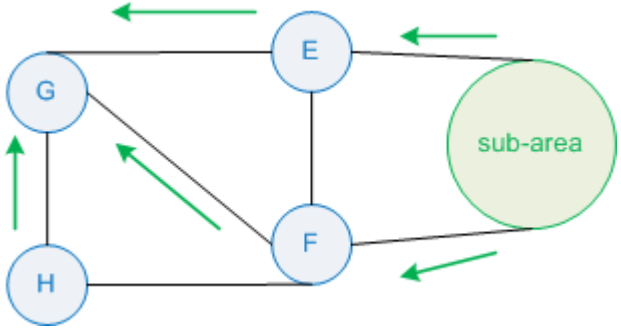

Finally, compute the destination-rooted MRTs. This shows the MRTs rooted at the proxy that represents E, F, G and H.
Note: A, B nodes sent non-MRT encapsulated packets to E,F.

Same method with proxy-nodes as for multi-homed prefixes.

IGP Extensions

- Future Extensibility built in
  - MRT Island Decision ID
  - MRT Algorithm ID
- Information to Share
  - Blue MRT MT-ID, Red MRT MT-ID
  - GADAG Root Election Priority
  - Red MRT Loopback Address, Blue MRT Loopback Address
- Forwarding Mechanisms Supported
- Capabilities Required & Available
  - IP FRR, LDP FRR, PIM FRR, mLDP FRR, PM Global Protection, mLDP Global Protection

Summary

- Filling in the details for a full solution
- Have modeled and verified correct alternate selection inside an area.
- Detailed feedback would be welcome.