PIM Extensions for Maximally Redundant Trees (MRTs)

Robert Kebler (Juniper)
Alia Atlas (Juniper)
Naiming Shen (Cisco)
Yiqun Cai (Microsoft)

PIM WG, Berlin, Germany

Background

- Growing interest in protecting Multicast traffic without requires Traffic Engineering.
- Live-Live solutions exist today but are dependent on the topology.
  - Idea is simple to easy to understand
  - Dual plane topologies with ECMP, LFAs
- New algorithms have been developed that can be used for Multicast protection.
  - Maximally Redundant Trees (MRT)
  - All topologies are supported
    - that have the physical connectivity (2-connected)
  - 100% coverage gives protection always, not just until the first network event
  - Need for protection should not dictate network topology

What are Maximally Redundant Trees

- We can always compute two link and node disjoint paths between any two routers?

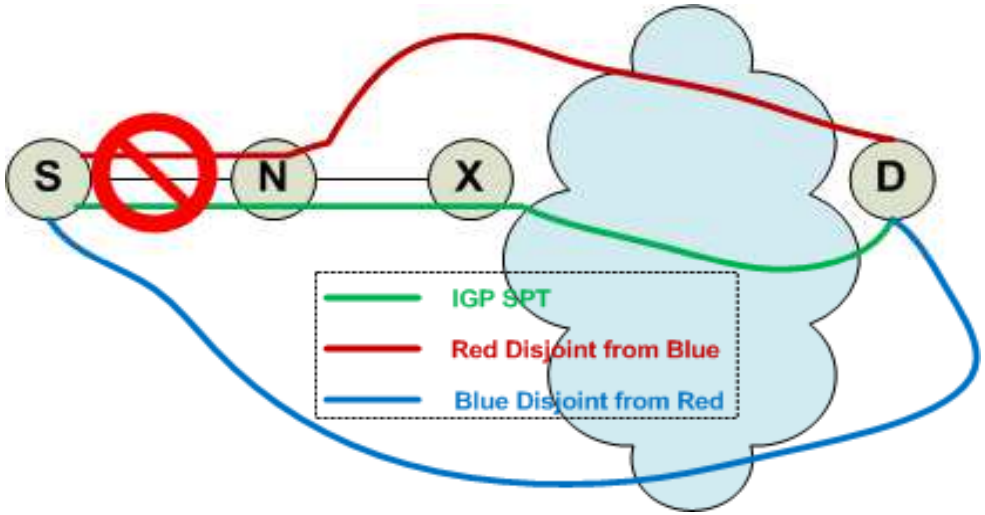

- The primary neighbor (N) **may** be on **at most** one of the red path or the blue path from S to D – but guaranteed not both.
- Algorithms give maximally disjoint paths
  - Handles all Network topologies

Types of Protection

- Local Protection (aka Fast-Reroute)
  - Point of Local Repair (PLR) **pre-computes an alternate** to avoid each local failure.
  - PLR sends traffic on alternate to Merge Point (MP)
  - Very fast repair time possible (< 50ms)
  - Examples:
    - RSVP-TE Fast-Reroute (RFC 4090)
    - IP/LDP Loop-Free Alternates (RFC 5286)
- Global Protection
  - Live-Live: Send traffic on two diverse routes and receiver/destination decides which to keep

Multicast Live-Live using MRT

- MRT is a natural fit for Multicast Live-Live
  - MRT algorithm provides two disjoint trees to reach the source
  - Last Hop router joins both trees and forwards a single stream to receivers
  - Adapts to topology changes
- Identify the MT-ID (Blue MRT or Red MRT) in PIM Join
  - Signaled with the MT-ID Join Attribute (RFC 6420)
- Stream Identification
  - Traffic can be distinguished on common link because of different Group (or Source).
  - Receivers join both the (S,G-blue) on the Blue MRT and (S,G-red) on the Red MRT.

Local Protection (Fast Reroute)

- Link-Protection:
  - PLR replication into tunnels using unicast route to reach next-hops
  - PLR learns of encap info from MPs in Join Attribute
- Node-Protection:
  - MP sends encap info in Join Attribute to its upstream, the Protect-Node
  - Protect-Node sends encap info for all MPs in Join Attribute
  - PLR must perform replication for all next-hops and next-next-hops

Local Protection Forwarding

- PLR may not be able to tell if a failure is link or node
  - Unlike unicast traffic where only getting to the destination matters, the next-hop router F cannot be bypassed because it might have local receivers.
  - Potential Merge Points (MP) must be able to handle receiving multiple streams of traffic and decide which to dump and which to forward.
  - A router only forwards alternate traffic when its upstream primary link(s) are down.
- PLR will send for a configurable amount of time to allow recovery
- Downstream routers should implement make-before-break

PIM extensions

- Changes to PIM:
  - Hello Message: Learn Capabilities
  - Join Attribute for PLR-driven tunnels
    - Send info for the MP

MRT Architecture

- MRT work being done in rtgwg
  - Multicast Architecture:
    - draft-atlas-rtgwg-mrt-mc-arch-02