83rd IETF – Paris, France

draft-wijnands-mpls-mldp-node-protection-00

IJ. Wijnands	ice@cisco.com
E. Rosen	erosen@cisco.com
K. Raza	skraza@cisco.com
J. Tantsura	jeff.tantsura@ericsson.com
A. Atlas	akatlas@juniper.net

Problem statement

- This draft documents a solution for mLDP node protection using unicast MPLS Tunnels
- Tunnels can either be a RSVP-TE P2P, LDP, LDP LFA, or something else
- Tunnels bypass the protected node
- mLDP packets get the Tunnel label pushed
- Need support for P2MP and MP2MP LSPs

Terminology

- Node protection using P2P Tunnels is all about the **PLR** learning the Merge Point (**MPT**) (leafs) of the protected node **N**.
- The PLR uses unicast P2P tunnels to bypass node N directly to the MPTs

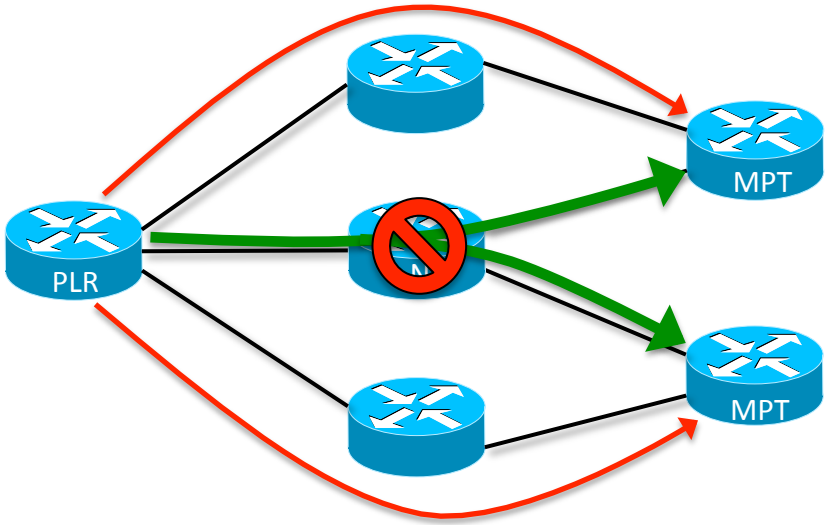

Solution

- In order to make this solution work, the PLR has to learn the remote bindings of N, called the Merge Points (MPT)
- Two solutions documented;
  1. This draft: based on Targeted LDP
  2. draft-zhao-mpls-mldp-protections

Solution – 1 (this draft)

- N advertises the PLR to its MPTs
- MPTs setup a T-LDP session with the PLR and advertise the bindings directly

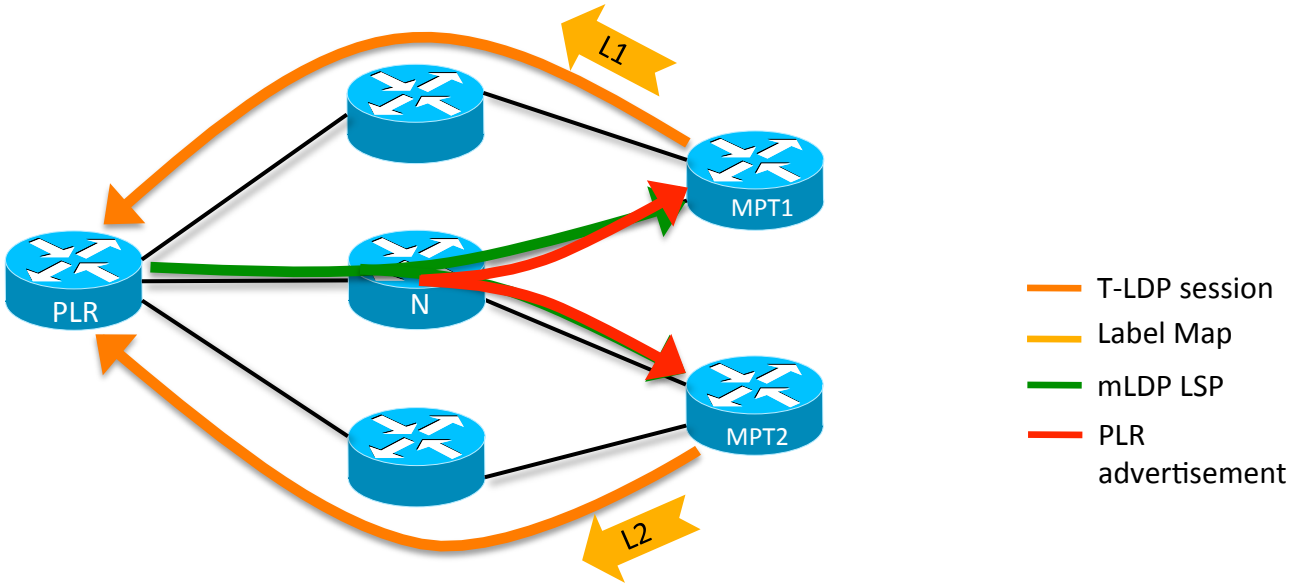

Solution – 1 (this draft)

- When node N fails, the T-LDP session remains up between the MPT and PLR
- Label Withdraw/Release messages can be exchanged
- The MPTs appear as regular bindings in the PLR forwarding table
- No special exceptions have to be defined to support MBB, GR, MP2MP, etc…
- This is at the expense of T-LDP session between the PLR and MPTs

Solution – 2 (draft-zhao)

- N advertises the MPTs to the PLR via its label mapping, as ships-in-the-night

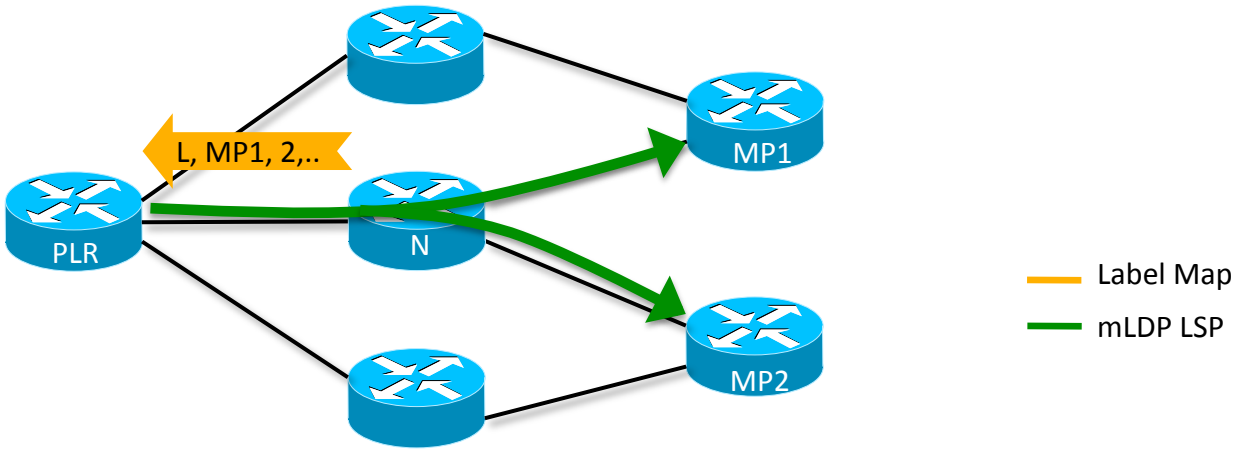

Solution – 2 (draft-zhao)

- It looks simple initially
- But as soon as N fails, the MPTs loose the signalling path to reach the PLR
  - Label Withdraw/Release messages can't be sent
  - Have to resort to timer based approach to withdraw and release labels
  - Potentially causes traffic gaps or duplication
- Support for MP2MP is not defined
  - Need 2-way path between PLR and MPTs

Solution – 2 (draft-zhao)

- Make-Before-Break is not defined
- Graceful Restart is not defined
- Typed Wildcard FEC is not defined
- Due to absence of LDP peering to exchange LDP messages, none of the existing features work by default.
- Many exceptions will have to be made in the code to support it.

Simplicity

- Using unicast MPLS tunnels is the simplest way to achieve Link and Node protection for Multicast.
- It piggy bags on the unicast infrastructure, mostly already in place.
- Downside is the replication load on the PLR.

Scalability related to T-LDP

- Targeted LDP sessions are per MPT – PLR pair, not per LSP
- Only need to exchange mLDP bindings
- If the fan-out on the PLR becomes a problem, maybe its better to not use P2P LSPs as backup but use P2MP LSPs
- Using P2MP has a totally different set of challenges being worked on currently

Conclusion

- The solution for mLDP node protection has to support the existing features as currently defined in the mLDP RFC 6388, like MP2MP, MBB, etc..
- Should not violate the LDP RFC 5036 due to not supporting Label Withdraw and Release for exchanged label bindings
- T-LDP is an architecturally clean way to address the problem. Don’t try to bypass it due to perceived scalability issues

Moving forward

- Working with the authors of draft-zhao-mpls-mldp-protections to resolve the difference of opinion.
- We are open to co-authoring