

T-LDP Hello Reduction

[draft-pdutta-mpls-tldp-hello-reduce-02](#)

Pranjal Dutta

Thomas Nadeau

Giles Heron

History

- 00-version was presented in IETF-74 in MPLS WG
- Addressed all feedbacks received from WG in revision -02 and presented in IETF-81 in PWE3.
- Moved from Informational to Standards Track as per suggestions from MPLS WG.

Objectives

- Optimize Targeted LDP (T-LDP) session maintenance in scaled T-LDP environment (e.g in mobility backhauls with Pseudowires/VPLS deployments).

Problem Statement

- After LDP session establishment T-LDP hellos become redundant. Session Keepalives are sufficient to maintain T-LDP adjacency.
- In scaled T-LDP environment , T-LDP Hellos are overheads after the session establishment if BFD etc. are used for tracking neighbor liveness.

Problem Statement

- LDP Hellos can cause DoS attacks leading to losing of established adjacencies. It is important to avoid dependency of hellos as far as possible.

Procedure

- Two T-LDP peering LSRs starts exchanging T-LDP hellos with “configured” hello timeouts as part of extended discovery.
- After LDP session is established, each LSR starts sending Hello Timeout of value 0xFFFF, which is infinite as per RFC 5036.
- If both peering LSRs start advertising Hello Timeout of 0xFFFF, then negotiated timeout falls back to 0xFFFF.

Procedure

- After the LDP session fails, both LSRs start transmitting T-LDP Hellos with the “configured” timeout values again.

Benefits

- Procedure is very simple and uses existing provisions of RFC 5036 . **Backward compatible with LSRs not implementing the procedure.**

Next Steps

- WG Document?