Reoptimization of Point-to-Multipoint Traffic Engineering Loosely Routed LSPs

draft-tsaad-mpls-p2mp-loose-path-reopt-03

Author list:

Tarek Saad (tsaad@cisco.com) - Presenter

Rakesh Gandhi (rgandhi@cisco.com)

Zafar Ali (zali@cisco.com)

Robert H. Venator (robert.h.venator.civ@mail.mil)

Yuji Kamite (y.kamite@ntt.com)

Outline

- **Scope and Requirements**
- **Problem Statement**
- **Signaling Extension**
- **Update and Next Steps**

Scope

- **P2MP-TE LSP [RFC4875]**
- **S2L Sub-LSP(s) signaled with Loose Hop ERO(s) or with no ERO [RFC3209]**
- **Loosely routed LSP reoptimization [RFC4736]**

Requirements

- As per P2MP-TE [RFC4875], an ingress node may:

1. Reoptimize the entire P2MP-TE LSP by resignaling all its S2L sub-LSP(s), i.e. all destinations, OR,
2. Reoptimize individual S2L sub-LSP, i.e. individual destination.

- [RFC4875] does not define mechanisms to reoptimize loosely routed (inter-domain) P2MP-TE LSPs.

Agenda

- Scope and Requirements
- Problem Statement
- Signaling Extension
- IETF Update and Next Steps

RFC4736 P2P LSP Reoptimization

Addresses reoptimization of loosely routed P2P LSPs

1. **Ingress sends "Path Re-evaluation Request" to trigger evaluation at midpoint LSR expanding loose next hops.**
   - **flag (0x20) in SESSION_ATTRIBUTES object in the Path message.**
2. **The midpoint LSR sends a (un)solicited "Preferable Path Exists" to notify the ingress node to trigger reoptimization.**
   - **PathErr code 25 (notify error defined in [RFC3209]) with sub-code 6.**

- **[RFC4736] does not define mechanism for P2MP-TE LSP Reoptimization.**

(Re-using) RFC4736 for P2MP-TE LSP Re-optimization

- Ingress sends "Path Re-evaluation Request" (PRR) for each individual sub-LSP to trigger evaluation at midpoint LSR expanding loose next hops
  - Ingress may have to send path re-evaluation requests on all (100s) sub-LSP(s) to decide whether or not to re-optimize the whole P2MP-TE LSP
  - Ingress may have to "heuristically" wait and aggregate all responses for "better path exists" to decide whether or not to do per sub-LSP or per LSP re-optimization
    - Ingress may prematurely start per sub-LSP re-optimization and then decide to abort and perform LSP re-optimization
    - Ingress may prematurely start re-optimization of sub-set of sub-LSPs, that may result in data traffic duplication [RFC4875] [Section 14.2]
  - May produce undesired results when inter-operating due to timing related issues and different implementations
- Can be avoided by extending the re-evaluation request messages for P2MP-TE LSP Tree reoptimization.

(Re-using) RFC4736 for P2MP-TE LSP Re-optimization

- Midpoint LSR sends an (un)solicited "Preferable Path Exists" (PPE) for each individual sub-LSP to notify the ingress node to trigger re-optimization
  - Midpoint LSR can not differentiate whether the request is to evaluate per sub-LSP path or whole P2MP-TE tree
    - May have to "heuristically" accumulate received requests for all sub-LSPs (using a wait timer) to interpret this as a re-evaluation request for the whole P2MP-TE LSP Tree
    - May prematurely notify better path exists for a sub-set of S2L sub-LSPs
  - Midpoint LSR may have to send better path exists on all (100s) sub-LSP(s) when it determine a better P2MP-TE tree exists
  - May produce undesired results when inter-operating due to timing related issues and different implementations
- Can be avoided by extending the notify messages send by the midpoint LSR for P2MP-TE LSP Tree reoptimization.

Agenda

- Scope and Requirements
- Problem Statement
- Signaling Extension
- IETF Update and Next Steps

Extensions For P2MP-TE LSP Tree Reoptimization

1. Ingress node sends “P2MP-TE Tree Re-evaluation Request" to query a a midpoint LSR for a preferable P2MP-TE LSP tree.
   - A new “P2MP-TE Tree Re-evaluation Request” flag is defined in Attributes Flags TLV of the LSP_ATTRIBUTES object [RFC5420] that is carried in a Path message
2. Midpoint LSR notifies ingress of solicited/unsolicited "Preferable P2MP-TE Tree Exists” node to trigger re-optimization of whole P2MP-TE LSP
   - Midpoint LSR sends a PathErr code 25 (notify error defined in [RFC3209]) with new sub-code "Preferable P2MP-TE Tree Exists”.
3. Any S2L sub-LSP of the LSP Tree transiting through the midpoint LSR can be selected to send the “P2MP-TE Tree Re-evaluation Request” to the midpoint LSR(s).
4. Notification of "Preferable P2MP-TE Tree Exists” can be sent back on the same S2L sub-LSP on which request was received on

Agenda

- Scope and Requirements
- Problem Statement
- Signaling Extension
- IETF Update and Next Steps

IETF Updates and Next Steps

- **Initial Draft was presented at IETF-89**
- **Draft was reviewed by Loa and MPLS-RT and comments were addressed by author(s) in version -03**
- **We would like to make this draft a WG Document.**

Thank You.