

draft-akiya-mpls-entropy-lsp-ping

IETF 88, Vancouver, Canada

Nobo Akiya
George Swallow
Carlos Pignataro
Nagendra Kumar

Background

- RFC4379
 - Defined an object for label load-balancing
 - Did not define usable procedures
- RFC6790 “Updated” RFC4379
 - But not with a complete description
 - In particular no way defined to put an entropy label in the Label Stack
 - Doesn’t cover PW or FAT PW cases
 - Doesn’t cover ELI/EL imposed along the path
 - Doesn’t cover hardware that can’t handle ELI/EL
 - Technically not an update at all by IETF procedures

Entropy Label FEC Stack Entry

- ELI does not need a new FEC, it can simply go in the NIL FEC
- Defined a new sub-TLV for EL
 - Value of label carried here
 - NIL FEC is not explicit enough in the PW, FAT PW cases

FAT PW

- Procedure for FAT PW is fairly simple
- Sender
 - Bottom of FEC Stack ::=EL
 - Next FEC pushed is PW FEC
 - Transport FEC(s) pushed
- Mid-point
 - Values of LB Type 9 successively replace value in EL

PW

- A bit less clean
- Trace using EL FEC instead of PW FEC
 - Advantage works exactly like FAT PW
 - Value of PW label goes into bottom of FEC Stack TLV and the actual FEC stack of the packet
 - Disadvantage – need to ping destination twice
 - Once to know you have arrived
 - Second time to verify PW
 - Alternative is to define new PW FEC types for when you have control words

LSP Ping/Trace on Entropy Motivation

- In general, 4 load balancer types
 - IP based LB, not imposing Entropy
 - IP based LB, imposing Entropy
 - Label based LB, not imposing Entropy
 - Label based LB, imposing Entropy
- Restore ECMP trace capability when mixed load balance techniques in
 - One LSP
 - Stitched LSPs
- Clarification on Multipath Type {9}

LSP Ping/Trace on Entropy

- Echo request
 - Multipath info to carry IP address set, Label set or both IP address and Label set.
- Echo reply
 - Communicate LB technique(s) to initiator
 - And one of below in multipath info:
 - IP address subset
 - IP address subset and corresponding entropy labels
 - Label subset
 - Label subset and corresponding entropy labels
 - None due to lack of multipath info match

LSP Ping/Trace on Entropy Extensions

- DS Flags:
 - L: 1=Label based LB, 0=Non-label based LB (i.e. IP)
 - E: 1=Imposes EL, 0=Does not impose EL
- Multipath Information Type {10}
 - IP Multipath Information
 - Label Multipath Information
 - Associated Label Multipath Information
- FEC definition for Entropy label
- Procedures on how to use above for initiator and responder nodes

Next Steps

- Work in progress
- The more we delved into it the more we found
- Would like the WG to review – more eyes, perspectives will help

Backup Slide

(in case we need to discuss with topology)

