Entropy Labels in MPLS Forwarding
draft-kompella-mpls-entropy-label-02

Kireeti Kompella, John Drake – Juniper Networks

Shane Amante – Level 3 Communications

Wim Hendrickx – Alcatel-Lucent

Lucy Yong – Huawei

History

- First presented back at IETF 73
- IETF 78: Update from -00 to -01
- Today (IETF 80): Update from -01 to -02

Background

- LAG and ECMP are powerful tools with widespread deployment
- Goal is to make them better
- Ingress PE's perform packet header key extraction and assigns an "entropy label" to incoming traffic
  - Typically, hash of 5-tuple of IPv4 or IPv6 header
- *Transit LSR's just use label stack (incl. entropy label) as input-keys* for LAG and ECMP, regardless of MPLS payload carried!
- Egress PE's discard entropy label before forwarding packet to final destination
- This draft is *complementary* to draft-ietf-pwe3-fat-pw

MPLS Label Stack

When app. labels ARE NOT used

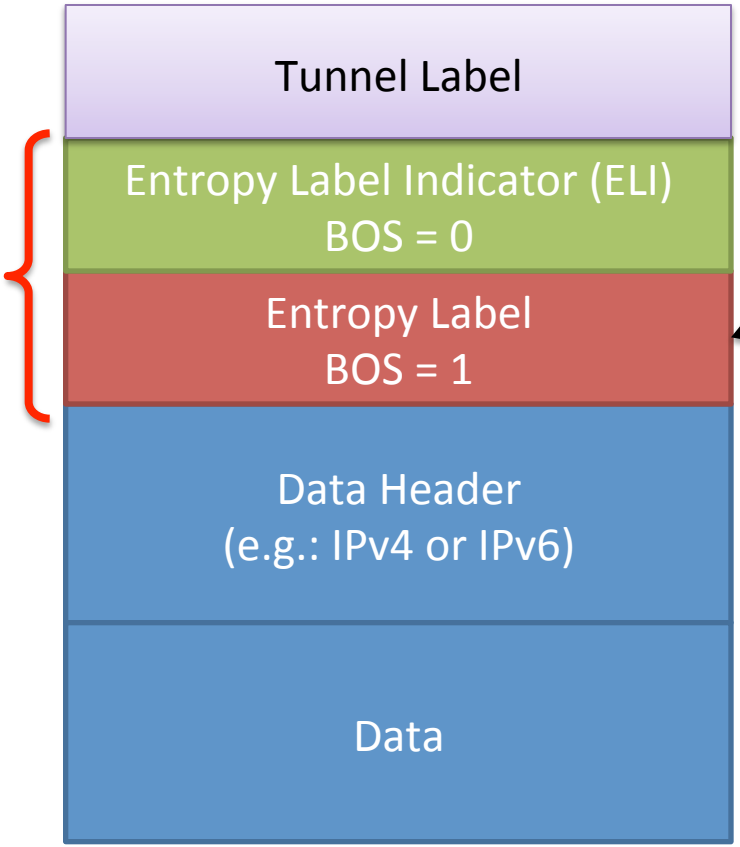

Entropy Label value

When app. labels ARE used

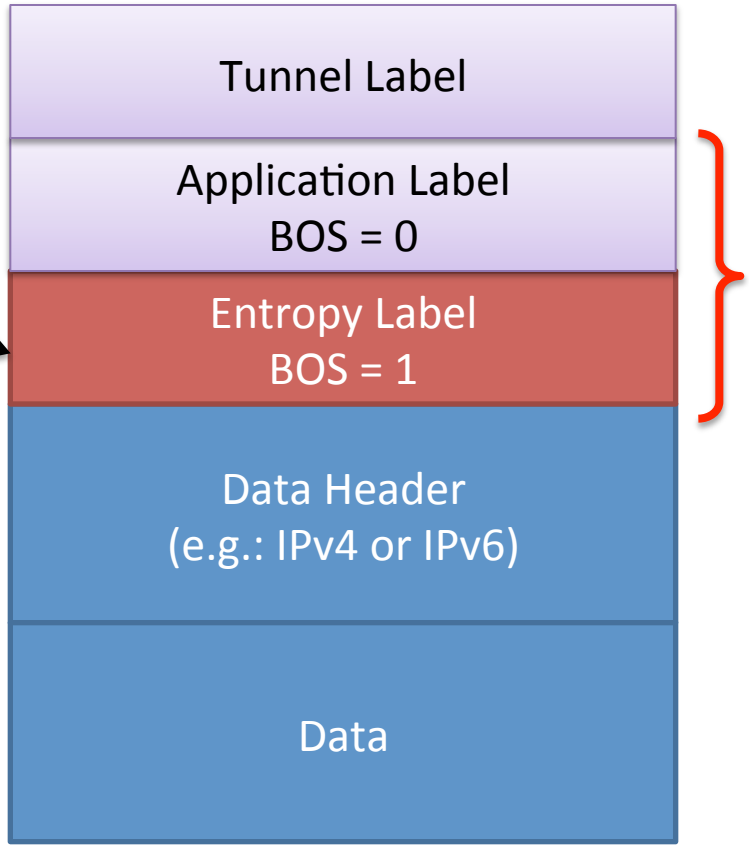

Changes from -01 to -02

- Substantial additions to doc
- **Section 5**: Signaling for Entropy Labels now includes BGP and, separately, RSVP-TE P2P (unicast, unidirectional & bidirectional) LSP Signaling Procedures
  - LDP Signaling for Entropy Labels already in draft
- **Section 6**: OAM + Entropy Labels
- **Section 7**: MPLS-TP + Entropy Labels
  - Entropy Labels do not apply to MPLS-TP LSP's, because MPLS-TP does not co-exist with ECMP per §3.2 of RFC 5921.
- **Section 8**: P2MP LSP's + Entropy Labels
- **Section 9**: (Example) Entropy Labels & Applications
  - Includes Tunnel LSP's, Inter-AS VPN's and Multiple Simultaneous Applications on same PE

Open Issues: 1 / 2
RSVP P2MP LSP Procedures

- **Problem**: receivers on same tree that:
  - Do and do not support receipt of Entropy Labels;
  - Join and leave P2MP LSP
- Proposed Solution:
  - Currently define that it's the responsibility of Ingress LSR to keep track of receiver's EL capabilities and determine when, and when not, to send EL's on the P2MP LSP.
  - Operators may decide to have two P2MP LSP's: first with P2MP EL capable receivers + second with P2MP non-EL capable receivers.

Open Issues: 2 / 2

OAM

- How will LSP ping/traceroute over LAG + ECMP work with entropy labels?
- Currently, LSP traceroute based on DPI of MPLS payload for input-keys, (e.g.: IP header 5-tuple), in real world
  - IOW, label range option of LSP traceroute not used in practice
- With EL's, ideally would have LSR's return label ranges to be used as input-key for load-balancing
- Problems:
  1. Not practical to globally (box-wide) disable DPI forLSP traceroute on LSR's, given EL capability or *lack thereof* on PE's
  2. How to coordinate, between PE's and LSR's, whether and how to enable/disable DPI vs. using a label range for LSP ping/ traceroute between PE's and transit LSR's

Next Steps

- Think the draft comprehensively addresses all signaling protocols and use cases
- Would like to ask to make this a WG draft, while continuing to resolve Open Issues