

# 82<sup>nd</sup> IETF - Taipei, Taiwan

**draft-wijnands-mpls-mldp-vpn-in-band-signaling-00**

IJ. Wijnands

ice@cisco.com

P. Hitchen

paul.hitchen@bt.com

N. Leymann

n.leymann@telekom.de

W. Henderickx

wim.henderickx@alcatel-lucent.com

# Problem we're solving

- [draft-ietf-mpls-mldp-in-band-signaling-04](#), documents a solution to map (1:1) Multicast Source trees (S,G) to mLDP LSPs.
- This is documented for the sources that are in the global table context, no VRF.
- The assumption at the time was that this is mostly used for IPTV like deployment which is in global table context.

# Problem we're solving (cont)

- Some providers use specific, dedicated VRF's to transport IPTV content across the network.
- The applicability of a simple 1:1 mapping of multicast Source trees to mLDP LSPs applies here as well.
- This draft documents the procedures needed to transport Multicast sources trees in a VRF context over mLDP LSPs via in-band signaling.

# Solution

- Multicast tree information, like (S,G) or (\* / M,G) is encoded in the opaque field of the mLDP FEC by the egress router(s).
- Including the RD of the Source (or RP) as advertised by BGP.
- The mLDP LSP is build through the MPLS core with (S,G, RD), or (\* / M,G, RD).
- This follows a similar approach as the VPN-recursive opaque encoding in:  
draft-ietf-mpls-mldp-recurs-fec-04


# Solution (cont)

- By adding the RD we're making the LSP unique in the core network.
- This prevents overlapping addresses.
- Each VPN IP multicast flow creates a unique LSP in the core.

# Opaque encodings

- We defined 4 new opaque encodings:
- Transit VPNv4 Source TLV
- Transit VPNv6 Source TLV
- Transit VPNv4 Bidir TLV
- Transit VPNv6 Bidir TLV

# Transit VPNv4 Source TLV



Type: (to be assigned by IANA).


Length: 16

Source: IPv4 multicast source address, 4 octets.

Group: IPv4 multicast group address, 4 octets.

RD: Route Distinguisher, 8 octets.

# Transit VPNv6 Source TLV



Type: (to be assigned by IANA).

Length:40


Source: IPv6 multicast source address, 16 octets.

Group: IPv6 multicast group address, 16 octets.

RD: Route Distinguisher, 8 octets.



# Transit VPNv4 bidir TLV



Type: (to be assigned by IANA).

Length:17


Mask Length: Length of the mask for group address, 1 octet.

RP: IPv4 multicast address, 4 octets.

Group: IPv4 multicast group address, 4 octets.

RD: Route Distinguisher, 8 octets.

# Transit VPNv6 bidir TLV



Type: (to be assigned by IANA).

Length:41

Mask Length: Length of the mask for group address, 1 octet.

RP: IPv6 multicast address, 16 octets.

Group: IPv6 multicast group address, 16 octets.

RD: Route Distinguisher, 8 octets.

# Scalability

- Each IP multicast flow maps to a LSP.
- Works well if the number of multicast flows is under control by the operator.
- mLDP is receiver driven, so its expected to scale well for thousands of LSPs.
- Scalability consideration is the same for both global table and VPN sourced multicast trees.

# Applicability

- This solution is NOT intended to be used as general purpose MVPN, like in; draft-ietf-l3vpn-2547bis-mcast-10
- It is a specific use-case where VRF' s just happen to be used for transporting IPTV like multicast traffic.
- Don' t panic.

# L3VPN

- This draft does describe minimal L3VPN specific procedures, the use of RD.
- draft-ietf-mpls-mldp-recurs-fec-04, does exactly the same.
- I propose to follow the same path as done for recurs-fec, do the work in MPLS WG, potentially do an last-call in L3VPN WG.

# Going forward

- Authors like to get comments on the draft.
- The intention is to make this a WG document.

Questions?