

# draft-ietf-l2vpn-evpn-01.txt

Sajassi (Cisco), R. Aggarwal (Arktan), W.  
Henderickx (ALU), N. Bitar (Verizon), A. Issac  
(Bloomberg), J. Uttaro (ATT), F. Balus (ALU), R.  
Shekhar (Juniper), J. Drake (Juniper), S.  
Boutros (Cisco), K. Patel (Cisco), S. Salam  
(Cisco)

IETF 84, July 2012

Vancouver

# Status

---

- Many discussions among the co-authors on the subjects that was listed in the previous IETF meeting
- Concluded along the line of objectives set out during the last IETF meeting – e.g., to reduce # of options where possible and to improve readability of the draft

# Delta from rev0

---

- Added new section 7 to describe Ethernet Tag and its value setting for different types of services
- Added new section 9 to describe all important concepts upfront in one place for better readability (previously nits and bits of it were spread throughout the draft)
- Reduced number of options for Ether AD route
- Removed 'don't-care' split-horizon label
- Removed source quenching option
- Added additional text for active/standby mode

# New Section 7 – Ethernet Tag

---

- Describes the relationship between Ethernet Tag in EVPN routes, broadcast domain IDs (e.g., CE-VIDs), and EVI
  
- Describes the setting for Ethernet Tag value in E-VPN routes for different services:
  - VLAN-based service Interfaces
  
  - VLAN-bundle service interfaces
    - Port-based service interfaces
  
  - VLAN-aware bundle service interfaces
    - Port based VLAN aware service interfaces

# New Section 9 – Multi-homing Functions

---

- Multi-homed Ethernet Segment Auto Discovery
- Fast Convergence & Mass Withdraw
- Split Horizon
- Aliasing
- Designated Forwarder Election

# Reducing # of options for Ether AD route

Flavor	A	B	C	D	E	F
<b>RD</b>	VPN RD	VPN RD	Segment RD	Segment RD	Segment RD	VPN RD
<b>Ethernet Segment ID</b>	VALID	NULL	VALID	NULL	VALID	NULL
<b>Ethernet Tag ID</b>	VALID	NULL	NULL	NULL	NULL	VALID
<b>MPLS Label</b>	VALID	VALID	NULL	NULL	VALID	VALID
<b>RT</b>	Single	Single	Multiple (corresponding to all VPNs on Segment)	Multiple (corresponding to all VPN instances enabled on PE)	Multiple (corresponding to all VPNs on Segment)	Single
<b>ESI MPLS Label Extended Community</b>	Not used	Not used	Contains the SH Label	Contains the SH Label	Contains the SH Label	Not used
<b>Use</b>	Advertise forwarding label per (ESI, Tag) for MPLS-based disposition.	Advertise forwarding label per VPN for MAC-based disposition.	<ol style="list-style-type: none"> <li>1. Advertise SH Label for an Ethernet Segment.</li> <li>2. Mass Mac withdraw upon a ES link failure</li> </ol>	<ol style="list-style-type: none"> <li>1. Advertise the special 'Don't Care' SH Label for ingress replication w/o source quenching</li> <li>2. Keep MPLS label stack consistent specially w/ flow label</li> </ol>	<ol style="list-style-type: none"> <li>1. Advertise forwarding label per ESI for MPLS-based disposition with label stack.</li> <li>2. Advertise SH Label for an Ethernet Segment.</li> </ol>	Advertise forwarding label per Tag for MPLS-based disposition with label stack.

# Don't care split-horizon label

---

- Removed this label and the associated mode in the Ethernet AD route
- Since Entropy label can be preceded with a reserved label, there is no ambiguity in MPLS label stack and thus no need for don't care split-horizon label

# Source quenching

---

- This option is used with ingress replication such that the replicating PE will replicate one BUM packet per egress sites per egress PE
- This option would allow for ingress PE not to send a packet to a multi-homed site that is participating in
- This option is very inefficient and thus it is removed
- As the result, segment-id in mcast route is no longer needed

# Active/Standby mode

---

- Added a section to describe Active/standby mode
- Active/standby flag is “ESI MPLS Label Extended Community” is set by the multi-homed PE advertising this route
- Remote PEs upon receiving this route can deduce that the received MAC advertisements from these sites, are sent by the primary PE
- Upon primary PE failure, if the site is dual-homed, then the remote PEs, can simply set the adjacency for these MAC addresses to the backup PEs for faster convergence and reduced flooding.

# Next Step

---

- Authors think that the draft is in a good enough shape for WG last call