

Extensions to BGP Signaled Pseudowires to support Flow-Aware Transport Labels

draft-keyupate-l2vpn-fat-pw-bgp-01

Bin Wen – bin_wen@cable.comcast.com

Keyur Patel – keyupate@cisco.com

Sami Boutros – sboutros@cisco.com

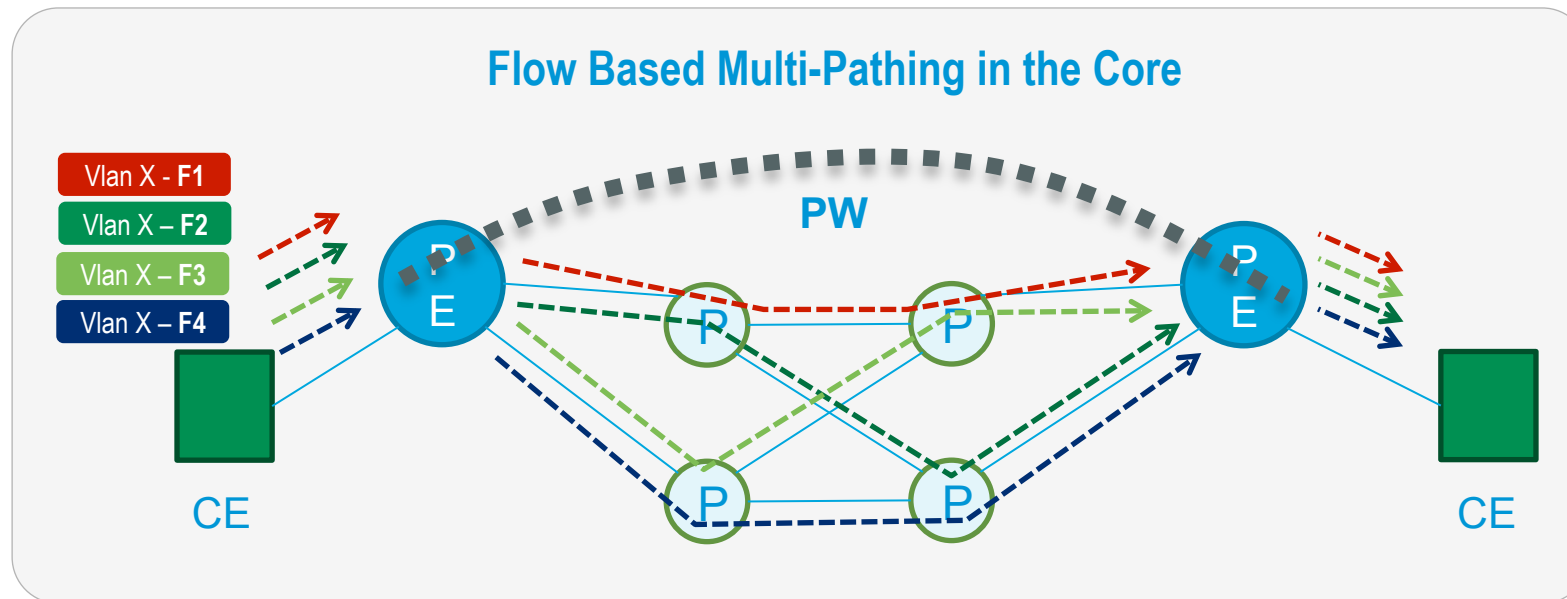
Jose Liste – jliste@cisco.com

Jorge Rabadan - jorge.rabadan@alcatel-lucent.com

IETF 90, July 2014

Problem Statement

- Ethernet services have become an important component of a SP product offering
- However, demand for high-speed Ethernet services (e.g. multi-GE or higher speeds) pose a problem for Network Operators as traffic from a given PW is not able to utilize all available paths (e.g. ECMP or LAGs) in the Core and instead it creates congestion in parts of the network
- Flow-based load-balancing in the Core becomes an important design consideration



Proposal

- This memo provides a solution for load-balancing of PW traffic with the following characteristics:
 - Based on **Flow Aware Transport PW** (IETF RFC 6391)
 - Applicable to deployments with **BGP-signaled VPLS (RFC4761)** and **BGP-signaled VPWS (RFC6624)**
 - Does **not require any forwarding behavior changes on transit LSRs**; i.e. NO changes to load-balancing hash functions on deployed P routers
- RFC4761 includes a Layer2 Info Extended Community in VPLS NLRI to convey information such as CW support, MTU, etc.
- **PROPOSAL – Use two (2) unused bits in Control Flag Bit vector to encode “T” and “R” bits as defined in RFC6391**

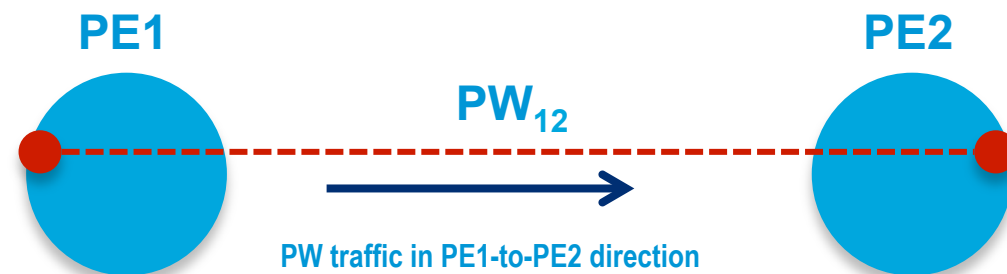


Figure 3: Layer2 Info Extended Community

FAT PW for BGP-sig VPWS / VPLS

T bit	Meaning
1	PE requesting to send PW traffic with Flow Label
0	PE will NOT send PW traffic with Flow Label

R bit	Meaning
1	PE willing to receive PW traffic with Flow Label
0	PE NOT willing to receive PW traffic with Flow Label

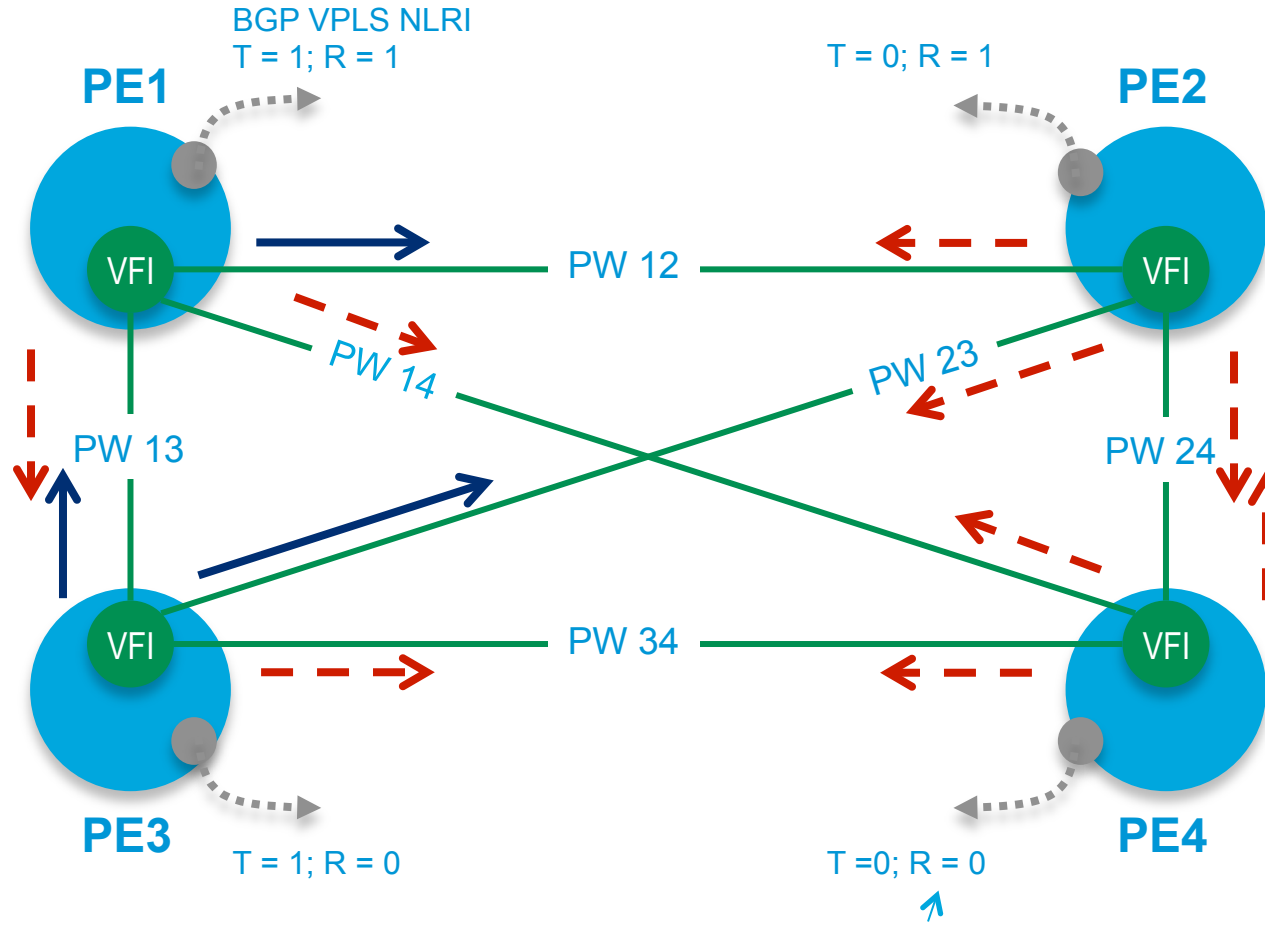


PE1 T bit	PE2 R bit	Meaning
0	0	FL NOT used for PW traffic in PE1-to-PE2 direction
0	1	FL NOT used for PW traffic in PE1-to-PE2 direction
1	0	FL NOT used for PW traffic in PE1-to-PE2 direction
1	1	FL used for PW traffic in PE1-to-PE2 direction

FAT PW for BGP-sig VPLS

PE	T	R
2	0	1
3	1	0
4	0	0

PE	T	R
1	1	1
2	0	1
4	0	0



PE	T	R
1	1	1
3	1	0
4	0	0

PE	T	R
1	1	1
2	0	1
3	1	0

Represents a "legacy" PE without FAT PW support

PW traffic with Flow label
 PW traffic without Flow label

Changes from version -00 to -01

- Modified title for better readability
- Added Jorge Rabadan (Alcatel-Lucent) as a new co-author
- Modified the location of T- and R-bits in the Control-Flag field of the Layer2 Info Extended Community (in order to prevent collision with other drafts)
- Added text to clarify the behavior for VPLS scenario where PEs may not share the same flow label settings
- Clarified text describing compatibility behaviors with PEs not supporting this draft
- Acknowledged contributions from John Drake (Juniper), John Brzozowski (Comcast) and Steven Cotter (Alcatel-Lucent)

Next Steps

- Authors believe that document is ready for WG adoption

THANK YOU !