

draft-singh-l2vpn-bgp-vpls-control-flags-00.txt

Updated processing of control flags for BGP VPLS

Ravi Singh (ravis@juniper.net)
Kireeti Kompella (kireeti@juniper.net)

IETF-88 (Vancouver)
speaker: Ravi Singh

BGP VPLS: Summary and an issue

- **Summary of BGP VPLS: efficiency of control messaging**
 - A PE's NLRI specifies behavior expected from all other PEs in that block (per VPLS):
 - VPLS label to use
 - Behavior regarding acting on control flags, etc
- **Issue:** PE behavior when mismatching settings of control flags:
 - Expectation of NLRI-originating PE cannot be met by every remote PE (w.r.t. acting on received control flags)
 - No support for selectively asking only some PEs to act on control flags
 - What do the other PEs do? What does NLRI-originating-PE do?
 - Ignore the mismatch?
 - Not bring up the PW?

What is this draft about?

- Supporting mixed-settings of control flags in a given VPLS:
 - Making BGP VPLS work when PEs differ in their settings of the control flags
 - For every (per-VPLS) PE-pair, if one PE does not have the same value for C-bit, the PW between the pair still comes up. However, does not use CW.
 - Similarly for other bits in the control-flags
 - Works in a backward compatible fashion
- Enabling use of p2mp LSPs as transport when PEs in per-VPLS PE-set have mismatching control flags

Illustrative BGP VPLS:

1/2

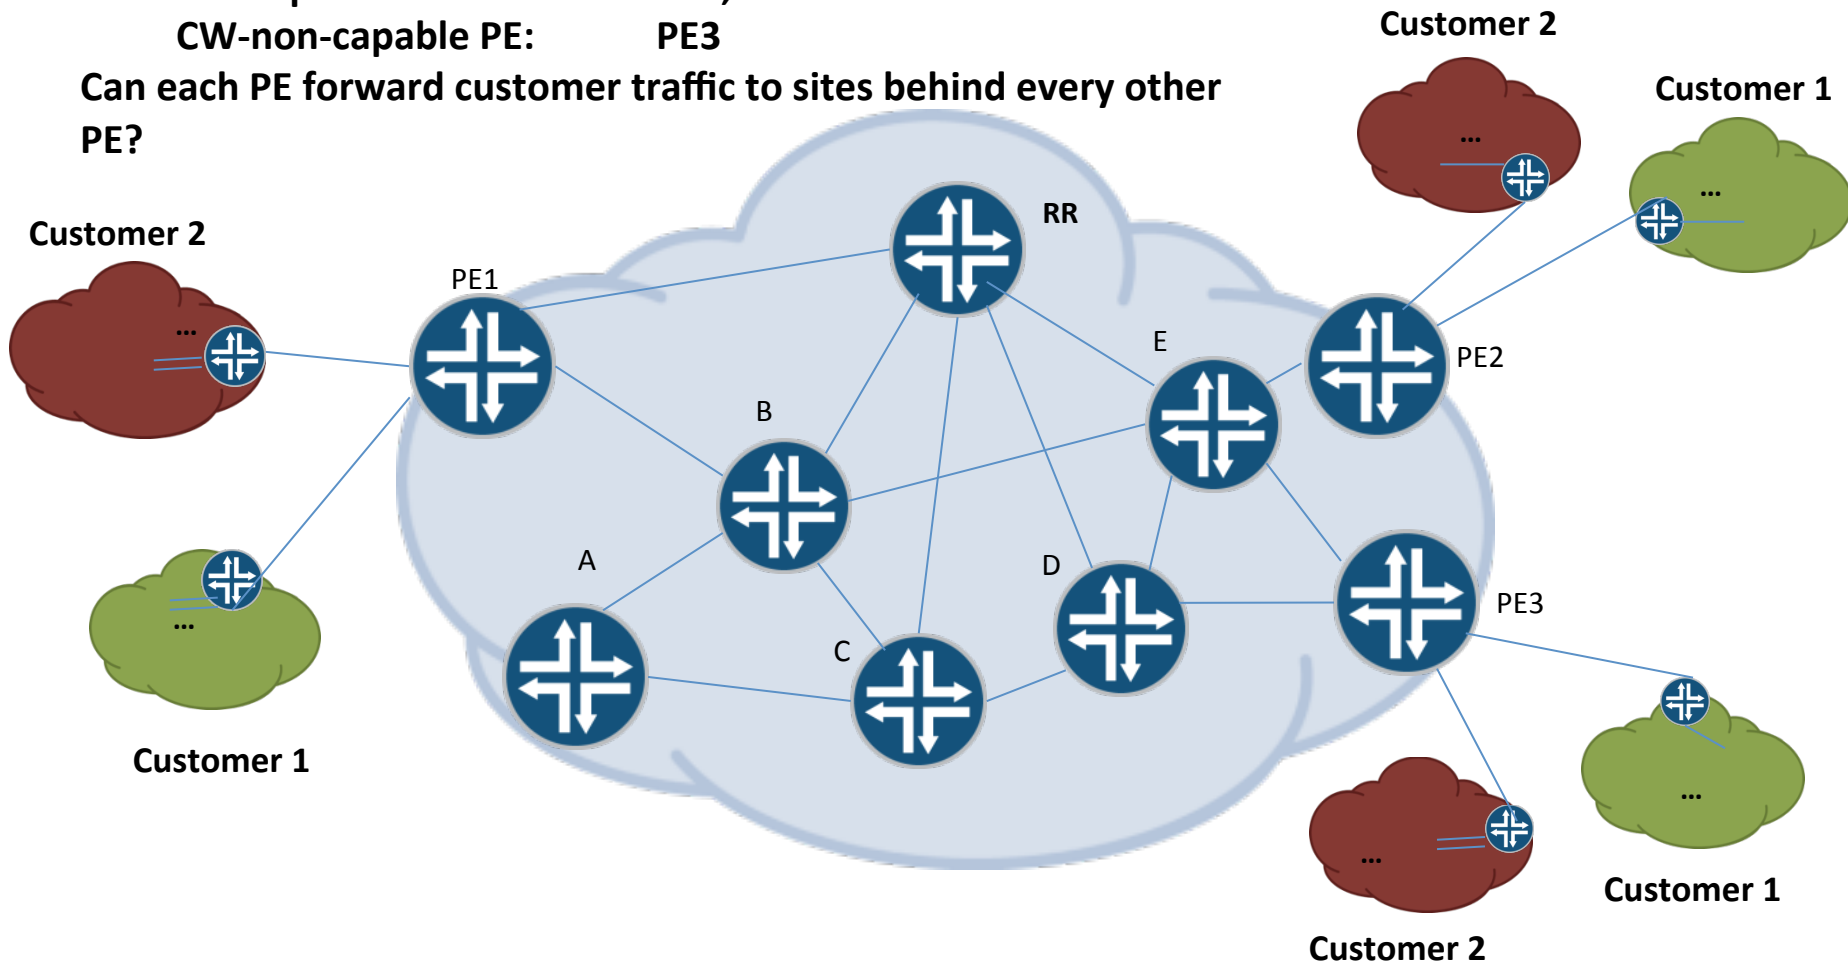
Using C-bit to illustrate the issue

Assume:

CW-capable PE s: PE1, PE2

CW-non-capable PE: PE3

Can each PE forward customer traffic to sites behind every other PE?



Illustrative BGP VPLS:

2/2

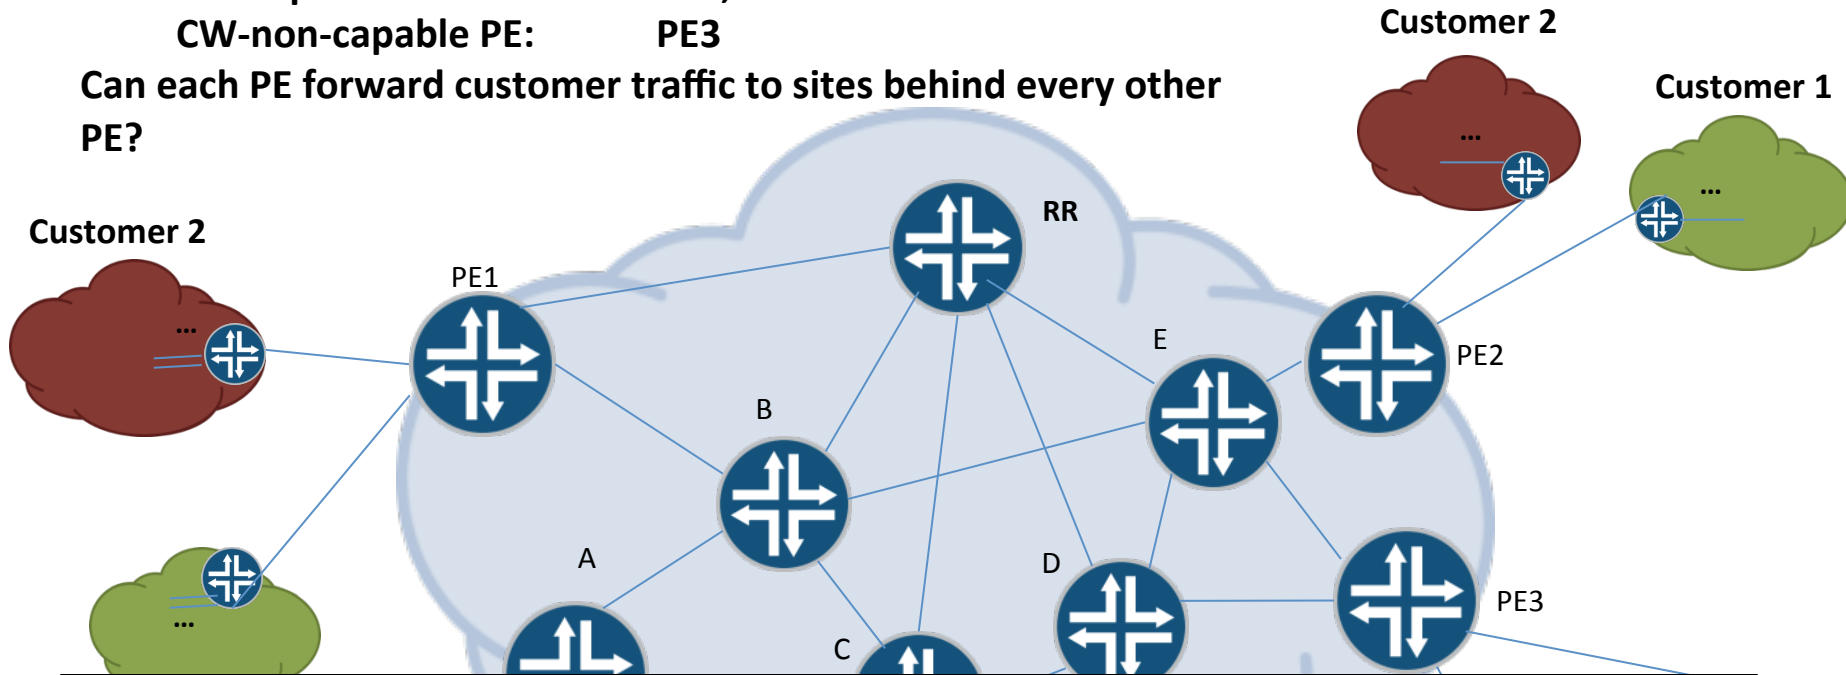
Using C-bit to illustrate the issue

Assume:

CW-capable PE s: PE1, PE2

CW-non-capable PE: PE3

Can each PE forward customer traffic to sites behind every other PE?



Before this draft:

UP PW (uses CW): PE1-PE2

Down PWs: PE1-PE3, PE2-PE3

With this draft: all PWs come up.

UP PW (uses CW): PE1-PE2

UP PWs (not using CW): PE1-PE3, PE2-PE3

Next steps

- Soliciting feedback
- Progressing in the WG