

BGP VPLS Multihoming

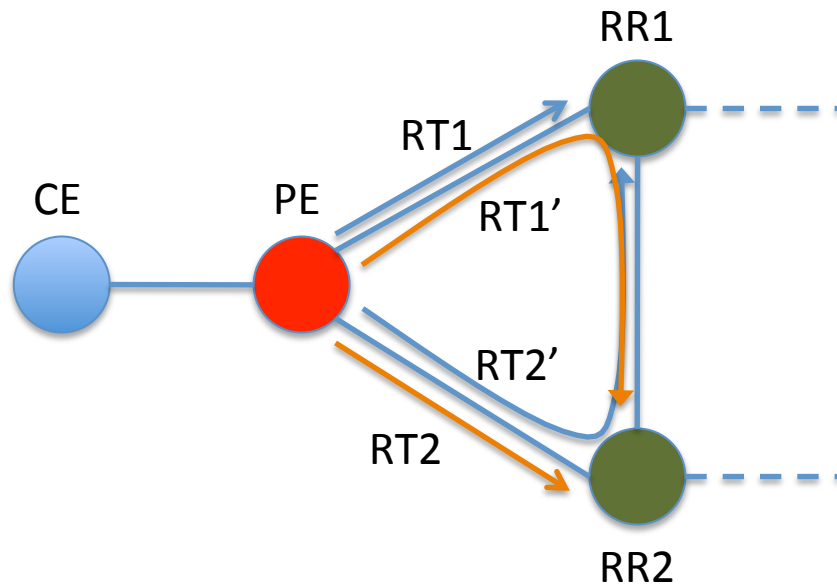
Draft-ietf-L2vpn-vpls-multihoming-04

- Bhupesh@Cohere.net
- Kireeti.Kompella@gmail.com
- Wim.Henderickx@Alcatel-Lucent.com
- Florin.Balus@Alcatel-Lucent.com
- Senad.Palislamovic@Alcatel-Lucent.com
- Wlin@Juniper.net
- Uttaro@Att.com

Problem statement

- Route Oscillation
- Architecture limitations due to not following standard BGP semantics
 - Algorithm defined
 - Standard BGP best path selection
 - VPLS DF Election
- Lots of not so clearly defined parameters throughout the entire draft
 - MH NLRI
 - Algorithm
 - VE-ID allocation
 - PW instantiation


Route Oscillation



The VPLS route for CE is advertised by PE to RR1 and RR2 as RT1 and RT2, respectively; RT1' and RT2' are these same routes re-advertised by RR1 to/from RR2

- RR1 receives RT1 from PE and RT2' from RR2. RT1 and RT2' are the same route (from PE). How should RR1 choose between them?
- RR2 has the same problem
- If RR1 chooses RT2' and RR2 chooses RT2, life gets difficult

eBGP vs. iBGP



If RR1 prefers RR2's route (iBGP) over RR3's route (eBGP), it will not advertise to its client PE1

Colored lines with arrows indicate VPLS route advertisement/re-advertisement

Proposed Solution

- Requiring unique RDs between MH PEs
- Elaborating on notion of bucketization and splitting BGP and VPLS algorithms fully
- Defining full BGP selection ensuring ebgp vs. ibgp and all other BGP rules
- Defining VPLS DF election rules

Some additional text clean up

- Made some portions of the text more readable for first time readers not familiar with BGP VPLS MH-ing concepts (there was a lot of assumptions)
 - As to why and how certain parameters and variables are used
- Backward compatibility issues
 - Defined PW Binding rules in relation to VE-ID allocation
- One more draft revision needed before final call