

VPLS Multi-homing


draft-ietf-l2vpn-vpls-multihoming-03

Kothari, Kompella, Hendrickx, Balus,
Uttaro

Draft has been stable

- ... so, time to shake things up :-)
- BGP DF election (more accurately, BGP route selection) as defined today has two problems:
 - Route oscillation
 - Longer than needed convergence and unnecessary traffic blackholing

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

Route Oscillation (2)

1. RR1 gets RT1 from PE. RR2 gets RT2 from PE
2. RR1 chooses RT1 as best path (only path)
 - RR2 chooses RT2 as best path
3. RR1 advertises RT1' to RR2
 - RR2 advertises RT2' to RR1
4. RR1 redoes route selection, picks RT2' instead of RT1 this time. RR2 similarly picks RT1'


Route Oscillation (3)

5. RR1 has a new best path, that of route RT2'; so it tries to re-advertise RT2' to RR2. But the best path is **from** RR2! Loop. So RR1 withdraws RT1'
 - Similarly, RR2 withdraws RT2'
6. With RT1' and RT2' withdrawn, we're back to step 1. Oscillate.

Long Convergence + Blackhole

- The current DF election rules say nothing about choosing between VPLS routes learned via eBGP and iBGP
 - Choosing an iBGP route means you are subject to the iBGP rules of not re-advertising that route to iBGP peers
 - Not getting that route may mean no backup path
 - This can adversely impact convergence time and cause a short blackholing of traffic

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

Solution

- Being wrangled, arm-wrestled and worked on in real time
- IDR chair suggests using standard BGP route selection in place of BGP DF election
 - This would simplify things protocol-wise, and bring back preference of eBGP over iBGP
 - However, there are issues with backward compatibility that need to be worked out

{VPLS | BGP} DF Election

- A philosophical issue stems from focusing on “DF election” versus “route selection”
 - DF election focuses on choosing among multi-homed sites to be the Designated Forwarder
 - However, there are times when one sees the same (single-homed) route from different peers with different paths and must choose among them
- A slight rewrite of the algorithms to reflect this would be helpful

Route Distinguishers

- RFC 4761 says that PEs supporting multi-homed sites MUST use (be configured with) the same Route Distinguisher (RD)
 - However, it is now common practice to use different RDs for multi-homed sites in IP VPNs
 - This practice gives BGP “add path” for free
- One more thing to think about: should this draft suggest (or require) this newer practice?