

Flow label for equal cost multipath routing in tunnels

draft-carpen-er-flow-ecmp-02

Brian Carpenter

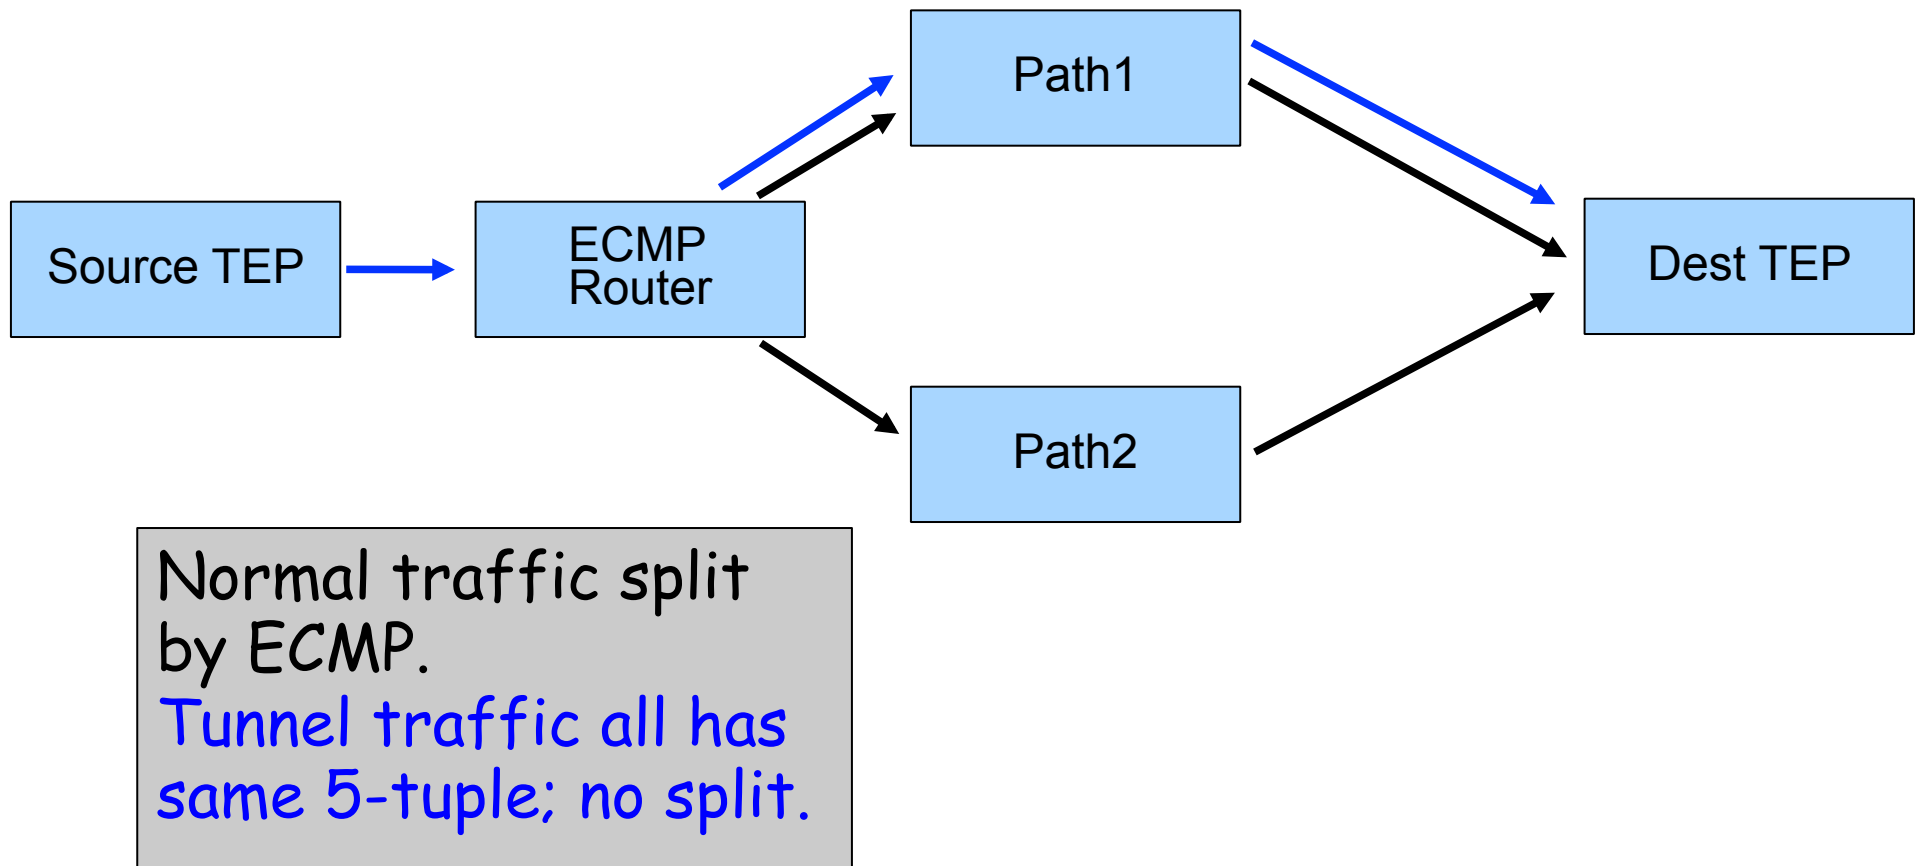
University of Auckland

Shane Amante

Level 3 Communications, LLC

July 2010

The problem with tunnels



Proposed solution

- For foo-in-IPv6 tunnels, the source TEP sets a flow label per user flow in the outer packet
 - For IP-in-IPv6, the flow label is based on the 5-tuple of the inner packet
 - It should be well distributed (pseudo-random)
- Intermediate ECMP or LAG paths use hash based on 6-tuple, the normal 5-tuple plus the flow label
 - works the same as before for non-tunnel traffic (and even better if flow label is set)
 - also splits tunnel traffic
 - fully conformant with RFC 3697
- Caveat: would not work in Inter-AS scenarios if the IPv6 flow label is allowed to be ‘mutable’ ...

Changes from -01 to -02

- Provided further background information
- Incorporated LAG (Link Aggregation)
- Changed to BCP language in Section 3, “Guidelines” to provide better guidance
- Added new co-author

Next Steps

- Adopt as WG draft?

Thank you!