draft-zzhang-bier-payload-label-00

BIER-WG, IETF124 Montreal

Jeffrey Zhang, Ed (zzhang@juniper.net),
Toerless Eckert, Ed. (Futurewei USA, tte@cs.fau.de)
Hooman Bidgoli (hooman.bidgoli@nokia.com)
Sandy Zhang (zhang.zheng@zte.com.cn)

Status

- No new version yet
  - but wanted to use the WG meeting time to discuss relevant discussion items
  - Did not get all players on the table. H-CA-Visa challenges – will be better @IETF125

- This is where we left @IETF123
  - Document proposes to replace semantic for code points 1,2 (MPLS over BIER)
    - From 1 = downstream, 2 = upstream
    - With 1 = CAN look up in global table, 2 = needs lookup in (per BFIR) context specific table (there is no MUST lookup in global table, but we will make that even more explicit!)
  - Still think this is logically and „extensibility wise“ the right approach:
  - Do not want to assign new code points for every new logical type of MPLS label
    - Today: Upstream, Downstream, SRGB, DCB
    - What comes next... ?
  - Code point is only relevant to the forwarding plane.
    - Do not put control plane semantic into it!
  - In the forwarding plane we only need to decide which LFIB we can do the lookup in

Do we even need 2 code points ? Why ? (Homan)

- We are a BFER
- BFIR-1 sends label L1 for VRF-VA, BFIR-2 sends label L1 for VRF-VB
  - Need to set up (BFIR-1, L1) -> VRF-VA and (BFIR-2, L1) -> VRF-VB
  - After parsing BIER header with packets from BFIR-1 or BFIR-2, we need to be able to do these context lookups
  - The BIER header code point can do that – but so far we only need 1.
- BFIR-3, BFIR-4 now also support DCB – but BFER not!
  - Configure („globally“) L1 to VRF-VA and L2 to VRF-B on BFIR-3, BFIR-4.
  - BFIR-1, BFIR-2 also do not support DCP
  - Assume BFIR-3 is on VRF-VA and BFIR-4 on VRF-VB – send according BGP of course
  - BFER establishes (BFIR-3, L1) -> VRF-VA, (BFIR-4, L2) -> VRF-B – ignores that they are signalled as DCB
  - Up to here – we still do not need 2 code points – all BIER/MPLS still looked up in context-specific LFIB.
- Now BFER is upgraded to also support DCB
  - Wants to do lookups in global LFIB to save not only labels – but total LFIB entries
  - ??? Potentially power performance doing context lookup ???
  - Now BFER would need a separate code-point for packets from BFIR-3, BFIR-4 to know that it can look up their packets in the global LFIB – as opposed to packet from BFIR-1, BFIR-2 (which it needs to lookup in context table)
- As soon as DCB is to be supported – seems we need two code-points
  - Assumption: We can not expect all BFIR to support DCB – because then again we would NOT need it.

Interop...

- Claim: DCB-only implementation is incompatible with non-DCB-only implementation
  - DCB-only implementation can only do lookup in global LFIB (else it could support non-DCB routes!)
  - non-DCB-only BFIR can send same label for different VRF (example: BFIR-1, BFIR-2).
  - Incompatible!
- HJN Interop test
  - N : non-DCB only. Simply ignored code-point, always did lookup in context LFIB
  - J: DCB only.
  - No failure ?!
  - Maybe just because there where no two non-DCB implementations sending the same label for different VPN
  - That would be the condition when the DCB-only implementation would have to fail
  - Fail could happen at control before traffic – recognizing the label collision
- Yes, No, Maybe... ??
- What do we call non-DCB ? SRGB-only ???

MTI – Mandatory to implement

- If you only want to implement one – DCB-only, non-DCB-only...
  - Which one should it be ???
  - Let‘s roll the golden dice: RFC decides the MTI
  - FCFS ? First RFC (SRGB) wins ???
- Or let the market decide
  - SPRING / RFC9800 – compressed SRv6 segment routing encap
  - Specifies two? (three ?) not sure ;-) incompatible options
  - Sounded terrible at first – but ultimately by now all implementations support all options (market decided)
- Proposal: No MTI in our doc („we‘re just doing what seems successful in unicast...“ ???)
  - But be very explicit about the interop issue for implementations implementing only one option

- What to say about possibly only partially working workaround like from interop ?
  - DCB-only implementation receiving non-DCB routes...
  - Can we make it work ?? Else it seems very risky

Operational Considerations

- It‘s the cool new thing – see draft-opsarea-rfc5706bis
- Would like to come up with text comparing SRGB vs. DCB approach to help operators
  - SRGB benefit: self-configuring, no controller needed to generate and distribute DCB config
  - DCB benefit: better scale. Roughly N times less LFIB entry requirements where N is average number of BFIR/VPN
  - Yes/No ?
- Anything else relevant ?

BGP

- How about the BGP signalling aspect...
- Tony mentioned we missed something where we may also need to update specs ?? Could not find it in the IETF123 notes...

- DCB config does not seem to be able to replace DCB signaling, because it does not say which BFIR is on which VPN...!!!
- If we would allow for that to be confired through DCB config – then we could forego the BGP DCB signalling...
- And what‘s N times more config sent if it‘s from a controller anyhow.... ??

The End