

BIER

Bit Indexed Explicit Replication Traffic Engineering

draft-eckert-bier-te-arch-03

IETF BIER-WG Buenos Aires 04/2016

Toerless Eckert, eckert@cisco.com

Gregory Cauchie, GCAUCHIE@bouyguestelecom.fr

Wolfgang Braun, wolfgang.braun@uni-tuebingen.de

Michael Menth, menth@uni-tuebingen.de


BIER-TE work since Yokohama

- Update -02 -03
 - Typos
 - Enhanced FRR section
 - > Two new co-authors
- Discus about feedback during Yokohama

BIER-TE FRR


- “Simple” idea:
 1. When BFR sees failure on link L (in bitstring), it triggers FRR.
Same rule as eg: RSVP/SR-FRR: Local detection.
 2. For FRR, packet gets modified:
 - Add some Bits – that’s the equivalent of the “backup path”
 - Remove some bits – that bit where the backup path + primary path (bits) could result in duplicates or even loops
 - Add/Delete Bitset is per-L – eg: failing link
- Challenge:
 - If we want to do node-protection, we need to know which next-next-hops we want to send to.
 - Need Structure: “If L fails”, here is set of per next-nexthop bit Add/Delete bitset.
 - Can not simply make per-L Add/Delete bitset send to all next-nexthops because that would create duplicates.

Some example




- Consider link failure on L or node failure on 5.
BFR 4 needs to assume what has happened when it sees L fail.
(assume link vs. node failure).
If node failure, BFR 14/15 become unreachable.

Some example



- Consider link failure on L or node failure on 5.
 - BFR 4 needs to assume what has happened when it sees L fail. (assume link vs. node failure).
 - If node failure, BFR 14/15 become unreachable.
- If 4 assumes link failure...
 - Would have to use some remote forward engineered adjacency to get packet copy to 5.
 - Delete bit to 5 add bit to get remotely to 5.
 - Could not avoid duplicate magenta/green traffic: magenta copy to 6/7/... was made independent of copy to 4.

Some example



- Assume node failure of 5.

If Bitstring has Link 5/16:

BFR 4 needs to send copy over to BFR 16.


Necessary in this example.

If Bitstring has Link 5/18:

BFR 4 needs to send copy over to BFR 17

Aka: not in this example

Some example



- Assume node failure of 5.
 - If Bitstring has Link 5/16:
 - BFR 4 needs to send copy over to BFR 16.
 - Necessary in this example.
 - If Bitstring has Link 5/18:
 - BFR 4 needs to send copy over to BFR 17
 - Aka: not in this example
- Can not rely on using existing branch 1/6/7/8/9/10:
 - Copies where created from 1 independent of branch of tree to

Some FRR thoughts

- Topology example is showing “worse case”
Resulting traffic paths during FRR not better than with older FRR solutions.
- Would like feedback how important this standard FRR paradigm is
- How about using instead end-to-end path protection
Eg: live-live. Algorithms exist, BIER-TE makes path engineering simple
- Or changing the signaling paradigm.
Past solutions where proposing fast propagation of failure to points that could make better FRR decisions.
Eg: BFD from BFR 4 to BFR 1 (BFIR). BFIR could make immediately optimal new Tree decisions.

Subdomains, SI

- Summary, discuss Yokohama / afterwards:

BIER-TE -02 define use of subdomains/SI with goal:

- Make operator experience easy
- Maximum reuse of “overlay” solution between BIER / BIER-TE
 - Without new work to make it work for BIER-TE

BIER arch author concerned about applying SI to TE:

- SI BIER specific part of automatic assignment of BFR-ID
 - And resulting forwarding in BIER network
 - Sequential assignemnt of BFR-ID in randomn places in network possible
 - BIER-TE:
 - Not all bits can be assigned to BFER
 - Assignment of Bits needs to take location into account
 - “need to have all bits for one region in one bitstring, not just “randomn”.

Conclusion

- New terminology for TE, no SI, improve comparison/differences explanation
- Analyse if different/improved label binding for BIER-TE required/beneficial.

Subdomains, SI

- Target Operator / Overlay experience:

Single instance of an overlay service (eg: L3VPN) is bound to one subdomain

So any overlay signaling that is signaling subdomains can get away just signaling the subdomain. And operator can just “configure” subdomain

- Overlay service uses sub-part of possible BFR and links
- If total number of links + BFR > Bitstring length: would need multiple subdomains. To stick with one subdomain / make operator life easy::
 - Subdomains has “sub-subdomains”. Propose to call them “slices”
 - Each slice = 1 bitstring within the subdomain.
 - BIER-TE SDN controller needs to help allocate bits for BFR and links into the slices to minimum number of slices is needed and replication efficiency is best.
 - Labels are bound to (subdomain,slice)
- Functionally, above is same as what was intended with -02/03 text, but want to avoid implying that Slices inherit any features that Sis have, but which are unique to BIER (see previous slide).



Questions ?

Improved
recovery

will get you a
new glass
when yours
breaks.