

Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP

draft-ietf-tsvwg-ecn-encap-guidelines-07

draft-briscoe-tsvwg-rfc6040bis-01

Bob Briscoe (Simula Research Lab)

John Kaippallimalil (Huawei)

Pat Thaler (Broadcom)

IETF-96, Jul 2016, Berlin

Bob Briscoe's contribution was part-funded by the European Community under its Seventh Framework Programme through the Reducing Internet Transport Latency (RITE) project (ICT-317700). The views expressed here are solely those of the authors.

Recap (1/2)

draft-ietf-tsvwg-ecn-encap-guidelines-07

- Purpose of this BCP draft:
 - Guidelines on addition of explicit congestion notification (ECN) to protocols that encapsulate IP,
 - e.g. tunnels, lower layers

- Not straightforward

- cross-organisation, cross-WG

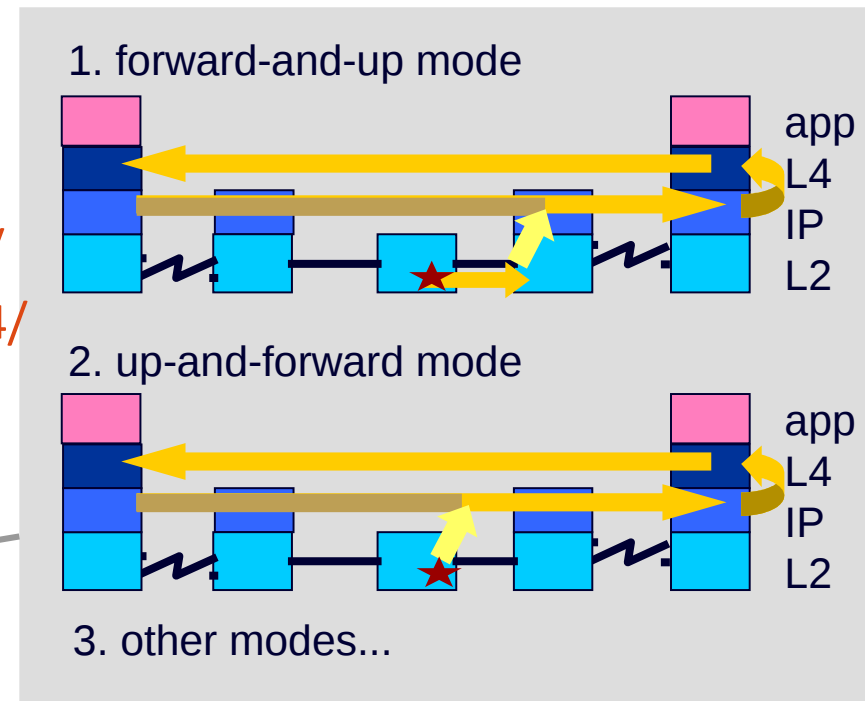
- IEEE: <https://datatracker.ietf.org/liaison/1364/>
- 3GPP: <https://datatracker.ietf.org/liaison/1424/>
- IETF: trill, nvo3, intarea (and previously mpls)

- cross-layer

- some lower layers have very different feedback structure

- incremental deployment

ECN propagation requires new logic in layer-egress and hosts




Recap (2/2)

Problem unique to ECN


- Both Diffserv (traffic class) and ECN have to propagate across layers

- DS propagates 'requirements' down



- ECN propagates...

- ECN-capable transport (ECT) down
- congestion experienced (CE) up




- ECN needs combination of inner and outer on decap

- see [RFC6040] for IP-in-IP

| incoming inner | incoming outer | | | |
|-----------------|----------------|---------|---------|------|
| | Not-ECT | ECT(0) | ECT(1) | CE |
| Not-ECT | Not-ECT | Not-ECT | Not-ECT | drop |
| ECT(0) | ECT(0) | ECT(0) | ECT(1) | CE |
| ECT(1) | ECT(1) | ECT(1) | ECT(1) | CE |
| CE | CE | CE | CE | CE |
| Outgoing header | | | | |

draft-briscoe-tsvwg-rfc6040bis* (1/2)


- Recently split out parts that update PS RFCs from draft-ietf-tsvwg-ecn-encap-guidelines (BCP)
 - likely to be fast-tracked
- Problem: RFC6040 “Tunnelling of ECN”
 - scope was only IP-in-IP tunnels
 - unclear whether this includes IP-shim-IP
- 6040bis solely extends scope of RFC6040
 - to include 'tightly coupled shim'
= shim added in same step as IP outer
 - “RFC 6040 SHOULD apply”
 - not MUST in case infeasible given structure of implementation



* Just an update, not a bis.
I didn't know that 'bis' is an IETF reserved word for a complete replacement.
If adopted, I'll use a different file-name.

draft-briscoe-tsvwg-rfc6040bis (2/2)

- rfc6040bis updates a number of PS tunnel specs (if approved)
 - RFC6040 ECN tunnelling (solely to widen scope)
 - RFC1701; RFC2784: GRE; RFC7637: NVGRE
 - RFC2661: L2TPv2; RFC3931: L2TPv3
 - RFC2637: PPTP
- Includes non-IETF specs with same structure that will need to be updated:
 - [GTPv1], [GTPv1-U], [GTPv2-C] GPRS Tunnelling Protocol (3GPP)
 - RFC7348: VXLAN
- aim:
 - if spec/implementation is being modified add RFC6040 support too
- rfc6040bis also lists specs that already require RFC6040 support
 - [draft-ietf-nvo3-gue] STD track Generic UDP Encapsulation
 - [draft-ietf-nvo3-geneve] STD track Geneve



Next steps

- **draft-ietf-tsvwg-ecn-encap-guidelines-07**
 - review from intarea / nvo3 please
 - comprehensibility? gaps?
- **draft-briscoe-tsvwg-rfc6040bis**
 - review from intarea / nvo3 please
 - is the list of tightly coupled shim specs complete?
 - would implementing RFC6040 with any of the listed tunnelling protocols present problems?