

# Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP (draft-ietf-tsvwg-ecn-encap-guidelines-07)

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

## 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-layer
    - some lower layers have very different feedback structure
    - incremental deployment
      - ECN propagation requires new logic in layer-egress and hosts
  - cross-organisation
    - IEEE: <https://datatracker.ietf.org/liaison/1364/>
    - 3GPP: <https://datatracker.ietf.org/liaison/1424/>

# recent activity

- 3GPP liaison – two outcomes:



- 3GPP specs mandate ECN marking where PDCP header has no ECN field
  - 3GPP-specific example added to Section 6 “Feed-Up-and-Forward Mode”

- Misunderstanding of AQM marking as on/off (for VoLTE only)

- Ought we to send a final formal liaison, giving pointers to advice?


- draft-ecn-encap-guidelines

- Split out updates to ECN tunnelling [RFC6040] into separate draft
  - intended, proposed standard – see next slide
- Closed off all open issues



# draft-briscoe-tsvwg-rfc6040bis\*

- Problem: RFC6040 “Tunnelling of ECN”; scope was all IP-in-IP tunnels
- 6040bis extends scope of RFC6040; to include 'tightly coupled shim' added in same step as IP outer
  - “RFC 6040 SHOULD apply”
  - not MUST in case infeasible given structure of implementation
- updates a number of PS tunnel specs (if approved)
  - RFC6040 ECN tunnelling
  - 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
- Also lists specs in progress that already cite RFC6040
  - [draft-ietf-nvo3-gue]      STD track      Generic UDP Encapsulation
  - [draft-ietf-nvo3-geneve]      STD track      Geneve



---

\* 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-ecn-encap-guidelines deltas between -05 & -07

- 1. Introduction: Added to list of examples of tightly coupled shims between IP headers
- 5.1. IP-in-IP Tunnels with Tightly Coupled Shim Headers
  - Replaced normative text with ref to new draft-briscoe-tsvwg-rfc6040bis
- 5.2. Wire Protocol Design: Indication of ECN Support: Added TRILL as an example of a well-design protocol that does not need an indication of ECN support in the wire protocol – see [draft-eastlake-trill-ecn-support]
- Encapsulation Guidelines: In the case of a Not-ECN-PDU with a CE outer, replaced SHOULD be dropped, with explanations of when SHOULD or MUST are appropriate
- Feed-Up-and-Forward Mode: Explained examples more carefully, referred to PDCP and cited UTRAN spec as well as E-UTRAN
- Added the people involved in liaisons to the acknowledgements
- Updated references
- Marked open issues as resolved, but did not delete Open Issues Appendix (yet)

# Next steps

- Finalise liaison with 3GPP?
- WGLC ecn-encap draft ... please
- draft-briscoe-tsvwg-rfc6040bis
  - adopt and fast-track? ... please
  - same as when it was in tunnel section of ecn-encap
  - just a different container