Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP

(draft-ietf-tsvwg-ecn-encap-guidelines-04)

Bob Briscoe (Simula Research Lab)
John Kaippallimalil (Huawei)
Pat Thaler (Broadcom)

IETF-95, Apr 2016, Buenos Aires

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-04

- Purpose of this draft:
  - Guidelines on standardisation of protocols and systems to ensure ECN is correctly deployed cross-layer

- Purpose of recent liaisons:
  - catch systemic ECN layering problems:
  - IEEE: https://datatracker.ietf.org/liaison/1364/
  - 3GPP: https://datatracker.ietf.org/liaison/1424/

Recent activity

- 3GPP response to IETF/tsvwg liaison identified
  - 6 3GPP WGs that could be affected by ECN layering
    - SA2, SA4, RAN2, CT1, CT3, CT4
  - 20 3GPP tech specs with normative refs or normative text relying on IETF RFC3168
  - We read them all…
  - Draft response prepared

SA	Service & System Aspects
RAN	Radio Access Network
CT	Core Network & Terminals

- Designing ECN support in TRILL (see later)

- Appropriate updates to draft-ietf-ecn-encap-guidelines-05
  - discussion on how to normatively update numerous tunnel specs

3GPP Liaison: Context

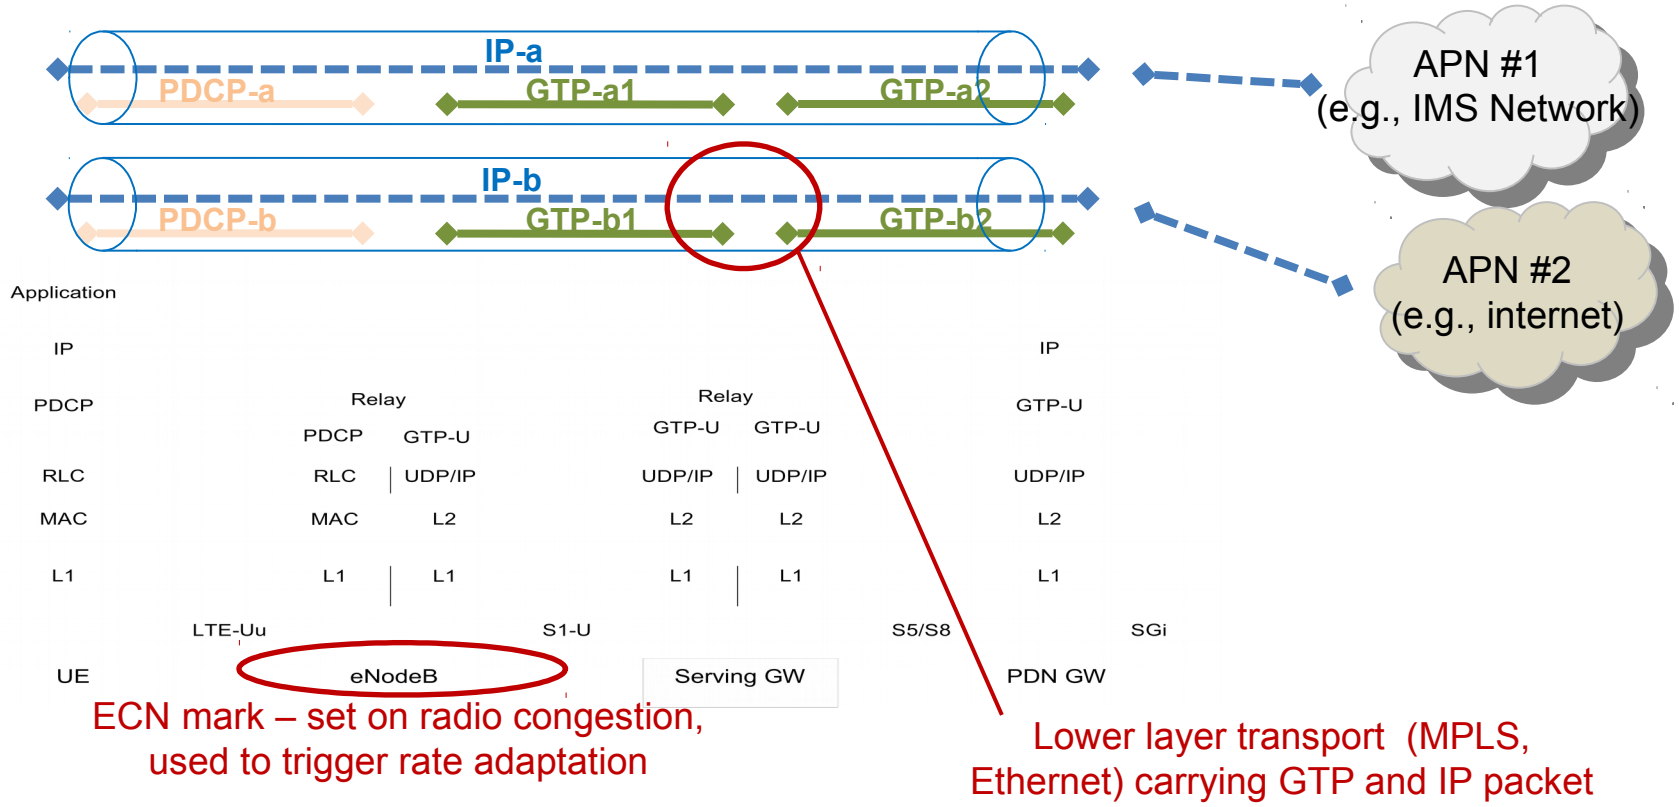

Snippets from liaison statement to 3GPP

" ….. However, ECN is now being used in a number of environments including coder selection and rate adaptation, where 3GPP protocols such as PDCP encapsulate IP. As active queue management (AQM) and ECN become widely deployed in 3GPP networks and interconnected IP networks, it could be incompatible with the standardized use of ECN across the end-to-end IP transport [draft-ietf-aqm-recommendation]. ……"

"The IETF is now considering new uses of ECN for low latency [draft-welzl-ecn-benefits] that would be applicable to 5G mobile flows. However, the IETF has realized that it has given little if any guidance on how to add explicit congestion notification to lower layer protocols or interfaces between lower layers and ECN in IP."

3GPP references to ECN

bold: LS says normative ECN text
italic: I found normative ECN text

- 3GPP TSG SA WG2
  - 23.228 IMS
  - 23.060 GPRS
  - ***23.401*** GPRS for E-UTRAN
- 3GPP TSG SA WG4
  - **26.114** IMS MM telephony; media handling & interaction
  - 26.223 Telepresence using IMS *[no ECN mention]*
  - 26.247 PS Streaming & 3GP-DASH *[no ECN mention]*
- 3GPP TSG CT WG1
  - 24.229 MM call control based on SIP and SDP
- 3GPP TSG CT WG3
  - **29.162** Interworking betw IP M-M (IM) Core Network (CN) subsystem & IP networks
  - 29.163 Interworking betw IM CN & Circuit Switched (CS) networks
  - 29.292 Interworking betw IM CN (IMS) & MSC Server for IMS Centralized Svcs (ICS)
  - 29.213 Policy and charging control signalling flows and QoS parameter mapping *[no ECN mention]*
- 3GPP TSG CT WG4
  - 23.333 M-M Resource Function Controller (MRFC) - Processor (MRFP) Mp i/f
  - **29.333** MRFC - MRFP Mp interface; Stage 3
  - ***23.334*** IMS Application Level G/w (IMS-ALG) - IMS Access G/w (IMS-AGW) I/f
  - **29.334** IMS-ALG - IMS-AGW; Iq Interface; Stage 3
  - **29.238** Interconnection Border Control Functions (IBCF) – Transition Gateway (TrGW) interface, Ix interface; Stage 3
  - **29.332** Media Gateway Ctrl Function (MGCF) - IM Media G/w; Mn i/f
  - **29.232** Media Gateway Controller (MGC) - Media Gateway (MGW) i/f; Stage 3
- 3GPP TSG RAN WG2
  - ***36.300*** E-UTRAN; Overall description
- 3GPP TSG RAN WG3
  - ***25.401*** UTRAN overall description

- Summary: all except the Radio Access ones are fine
  - primarily L4-7 and fully compatible with ECN-in-RTP [RFC6679]

ECN in 3GPP TS 25.401 & TS 36.300

Overall description UTRAN & E-UTRA, respectively

7.2.11 & 11.6 respectively: "Explicit Congestion Notification

The eNB **should** set the Congestion Experienced (CE) codepoint ('11') in PDCP SDUs in the downlink direction to indicate downlink (radio) congestion if those PDCP SDUs have one of the two ECN-Capable Transport (ECT) codepoints set. The eNB **should** set the Congestion Experienced (CE) codepoint ('11') in PDCP SDUs in the uplink direction to indicate uplink (radio) congestion if those PDCP SDUs have one of the two ECN-Capable Transport (ECT) codepoints set.

"

This gives normative specification of base station behaviour

UTRAN	Universal Terrestrial Radio Access Network
E-UTRA	Evolved UTRA
PDCP	Packet Data Convergence Protocol
SDU	Service Data Unit
eNB	evolved NodeB (base station)

- 2 problems (next 2 slides):
  1. PDCP layered below IP and has no ECN field of its own
  2. marking behaviour unclear
- The liaison has highlighted difficulty understanding RFC3168

#1 ECN layering problems in 3GPP TS 25.401 & TS 36.300

- no specs of how the ECT codepoint got to the outer header at the PDCP
  - GTP tunnel end-point specs need to refer to RFC6040 (bis?)
- Seems to solely apply to downlink
- eNB replaces IP-UDP-GTP header with RLC-PDCP which has no ECN field
- "ECN codepoints in PDCP SDUs" must mean "ECN in the IP header inside"? Otherwise:
  - how does eNB know a PDCP SDU has ECT codepoint set at the IP layer?
  - how does PDCP propagate ECN to the IP header at the terminal?
  - but the IP headers that PDCP encapsulates, may be compressed (RoHC)

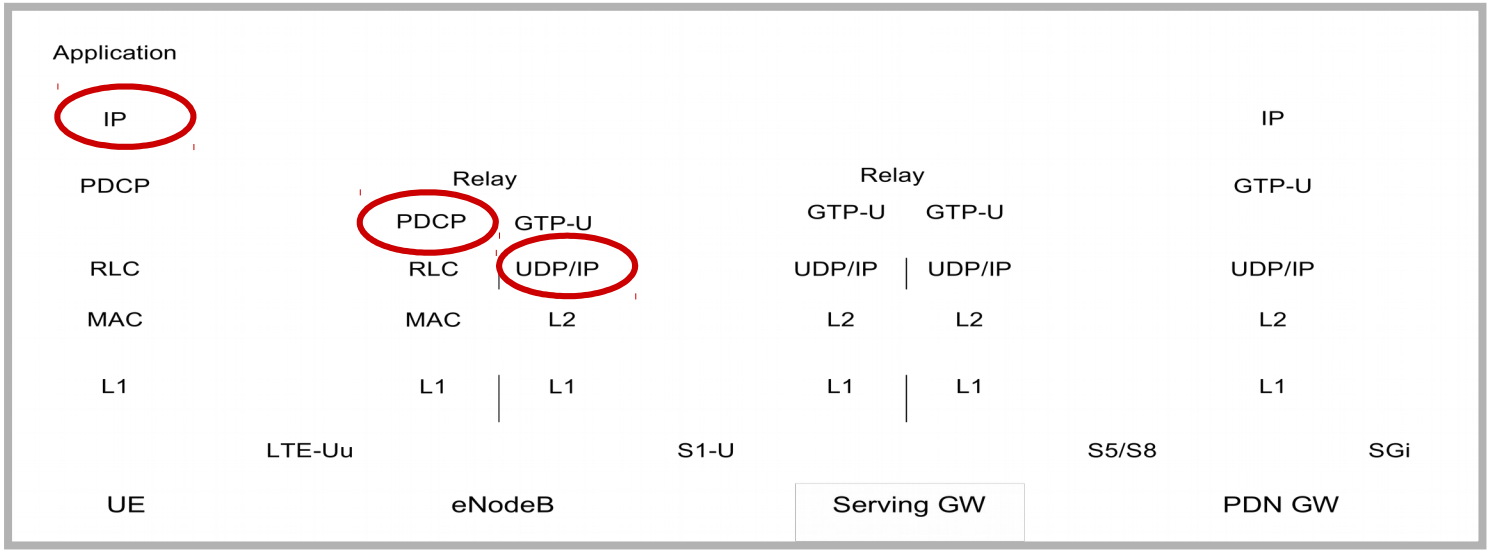

#2: ECN behaviour in 3GPP TS 25.401 & TS 36.300

- Marking & response behaviour – 3GPP needs to clarify:
  - whether marking is confined to voice bearers (and why?)
  - whether the wording implies “all or nothing” marking
  - otherwise incompatibility between “all-or-nothing” and “loss-equivalent” [RFC3168] marking in other networks
    - Should codec rate reduction be triggered on a single “CE” mark?
    - Should codec rate reduction be triggered on multiple “CE” marks?

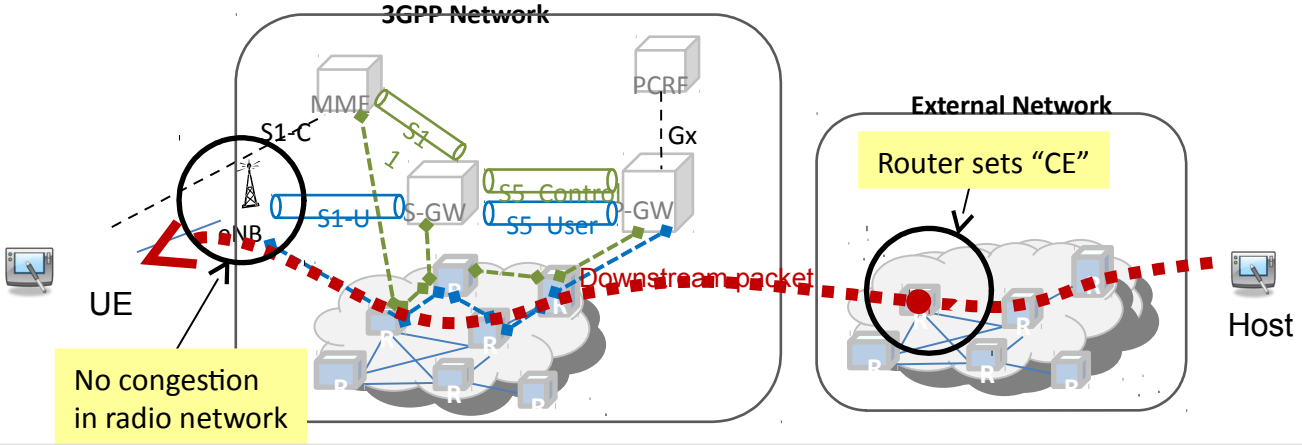

TS 23.401 specifies some constraints on the marking behaviour, but I could not find a spec of the behaviour itself:
“To make sufficient time available for end-to-end codec rate adaptation the E-UTRAN/UTRAN **should** attempt to not drop any packets on a bearer for a default grace period of at least 500 ms after it has indicated congestion with ECN on the bearer for packets within the packet delay budget. During this ECN grace period the E-UTRAN/UTRAN **should** also attempt to meet the QCI characteristics / QoS class associated with the bearer.
NOTE 1: Note that the receiving end-point should interpret all ECN-CE signals received within one end-to-end round-trip time as one "congestion event" (see IETF RFC 3168 [55] and TS 26.114 [56]).”

ECN Support in TRILL

- A new technique for adding ECN support at a lower layer
  - only if L2 protocol is well-designed for forward-compatibility (as TRILL is)
  - to be added to list of ECN encap tricks in ecn-encap-guidelines draft
- TRILL has a facility to flag a critical extension
  - that the egress must drop if it doesn't understand
  - exploit this for congestion experienced flag
  - legacy egress that does not understand how to propagate ECN drops the frame (desired outcome)
- New individual draft-eastlake-trill-ecn-support-00
  - but it currently does not cover this new idea
  - see TRILL proceedings this week for design slideware

Updating Tunnel Specs

- ecn-encap-guidelines-05 includes following text:

```
Therefore, in all such tightly coupled IP-in-IP tunnelling
protocols the rules in [RFC6040] for propagating the ECN field
between the two IP headers SHOULD be applied directly.

Examples of tightly coupled IP-in-IP tunnelling protocols where
[RFC6040] can be applied directly are:

    o L2TP [RFC2661]
    o GRE [RFC1701], [RFC2784]
    o PPTP [RFC2637]
    o GTP [GTPv1], [GTPv1-U], [GTPv2-C] {Note 1}
    o VXLAN [RFC7348].
```

- Given the intent is to update all these Proposed Standards, how best to do this?
  a) rephrase ecn-encap-guidelines to update RFC6040 and all these RFCs
  b) new 1-page RFC6040bis Proposed Standard to extend scope to these RFCs

{Note 1} GTP is a 3GPP protocol sandwiched within IP headers

- the IETF would specify ECN propagation across the IP headers
- 3GPP would refer to that in the GTP specs

Next Steps

- IETF response to 3GPP liaison response
- ecn-encap edits:
  - RFC6040bis text (either in ecn-encap, or new draft)
  - generalisation of new TRILL technique
- then WGLC ecn-encap-guidelines ?

Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP
(draft-ietf-tsvwg-ecn-encap-guidelines-04)

Q&A

Spare slide

#2: Ethernet Backhaul in 3GPP Networks

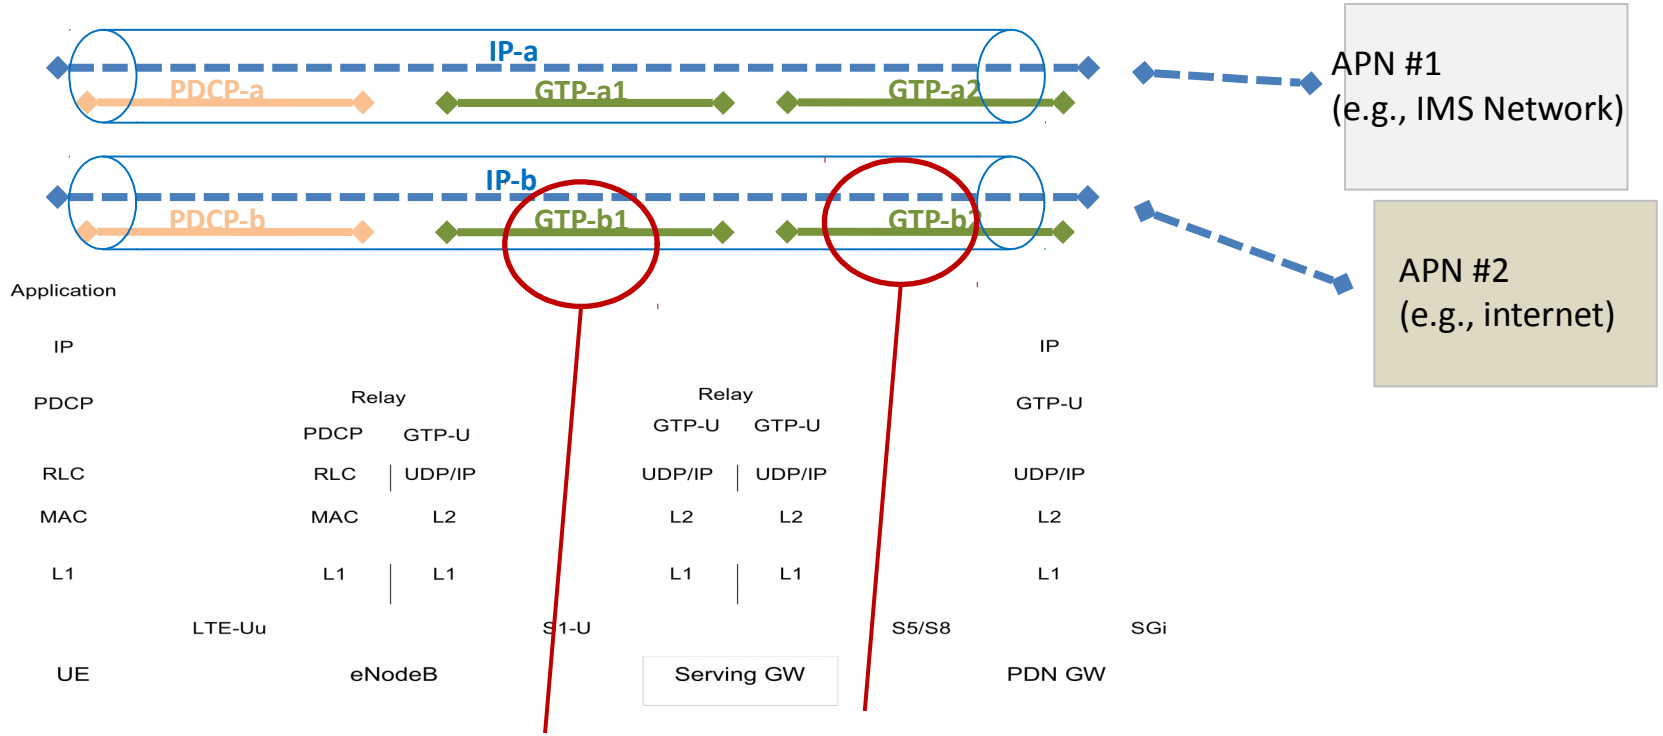

Ethernet backhaul for S1 and S5

Expected behavior:

- If congestion experienced, unlike in MPLS (feed-forward-and-up), the Ethernet backhaul network should set “CE” in IP (feed-up-and-forward)