Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP

draft-briscoe-tsvwg-ecn-encap-guidelines-00

Bob Briscoe

IETF-80 Mar 2011

This work is partly funded by Trilogy, a research project supported by the European Community www.trilogy-project.org

explicit congestion indications from lower layers

problem: standardise interface with IP

- switches can ‘mark’ Ethernet header
  - using AQM[1] developed for IP or MPLS
- ‘mark’ may change CoS[2] or a spare bit
  - but no Ethernet standard for this
- L2 congestion notification stds exist
  - typically limited to subnet
- pressure to link these subnets
  - using IP as portability layer
- lower layers need guidelines
  - to interface to ECN[3] in IP [RFC3168]

positioning within canonical congestion control architecture

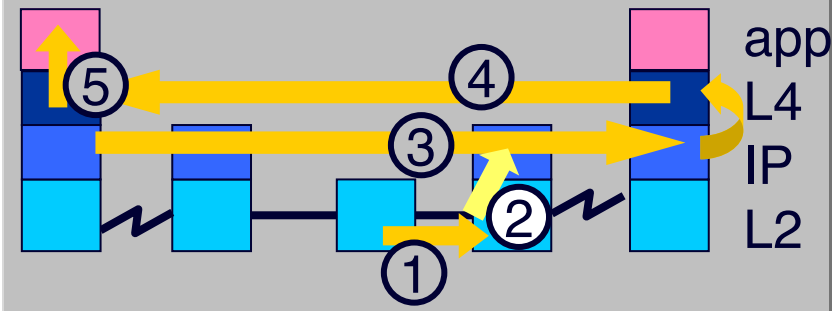

1) congestion indication in lower layer
2) propagate up to ECN in IP — focus
3) ECN in IP
4) feedback in e2e transport
5) response to congestion

AQM = active queue management (e.g. RED)
RED = random early detection
CoS = class of service in IEEE 802.1p
ECN = explicit congestion notification

involves messing with

the neck of the hourglass

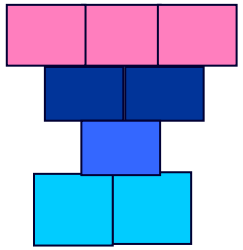

- avoid precluding L2 innovation
- must not be over-prescriptive

- wide review necessary

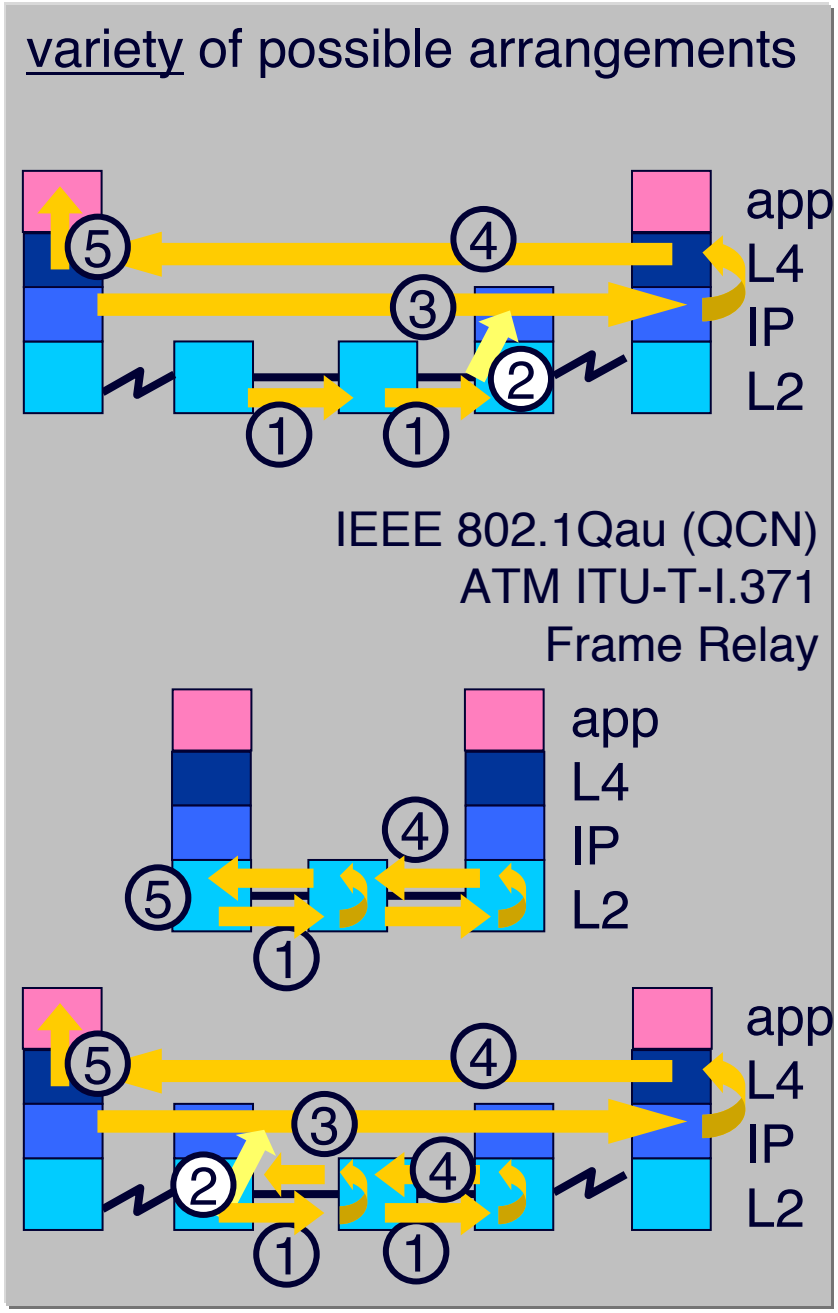

status of congestion notification
in protocols that encapsulate IP

- IETF

  done: MPLS-in-MPLS, IP-in-MPLS [RFC5129], IP-in-IP [RFC6040]

  to do: trill-rbridge-options (in progress),
  & pass ECN thru tunnel protocols, eg. L2TP, GRE

- Other standards bodies:

  done: QCN [802.1Qau], Frame Relay, ATM [I.371]
  (all subnet-local)

  todo: IEEE 802.1, (802.3, 802.11), …?
  & pass ECN thru tunnel protocols, eg. 3GPP GTP

L2TP = layer 2 tunnelling protocol [RFC2661]
GRE = generic routing encapsulation [RFC1701, RFC2784]
QCN = quantised congestion notification
GTP = GPRS tunnelling protocol [3GPP TS 29.060]

the main problem: incremental deployment

- IP-ECN designed for incremental deployment

		congested queue supports ECN?	
transport supports ECN?	IP header	N	Y
N	Not-ECT	drop	drop
Y	ECT	drop	CE

- if transport only understands drop
  - lower layer must not send it congestion indications
- need not mimic IP mechanism (grey)
  - but needs to achieve same outcome (white)

ECT = ECN-capable transport
CE = Congestion Experienced

guidelines

- identifying whether transport will understand ECN
- propagating ECN on encapsulation
- propagating ECN on decapsulation
- reframing issues

guidelines

- identifying whether transport will understand ECN
  - new problem: will decapsulator understand ECN?
- propagating ECN on encapsulation
  - copying ECN down for monitoring purposes
- propagating ECN on decapsulation
  - combining inner & outer
- reframing issues
  - marked bytes in ≈ marked bytes out
  - timeliness – don't hold back any remainder

next steps

- process
  - adopt as wg item?
  - will require liaison with other standards bodies
  - informational or best current practice?
- document
  - add architecture diagram(s)
  - want to avoid precluding L2 innovation – need help
  - it just mentions that L3 switches mark IP-ECN
    - doesn't say whether good or bad
    - I'd like to say it's OK: any objections?
  - to address: tunnelling protocols if never outer on the wire

Guidelines for Adding Congestion Notification to Protocols that Encapsulate IP

draft-briscoe-tsvwg-ecn-encap-guidelines-00

Q&A