making stuff real
re-feedback

Bob Briscoe, BT Research

Nov 2005

CRN DoS resistant Internet w-g

the problem: accountability for causing congestion

- main concern
  - non-compliance with e2e congestion control (e.g. TCP-friendly)?
  - how can ingress netwk detect whole path congestion? police cc?
- not just per-flow congestion response
  - **smaller:** per-packet
    - single datagram ‘flows’
  - **bigger:** per-user
    - a congestion metric so users can be held accountable
    - 24x7 heavy sources of congestion, DDoS from zombie hosts
  - **even bigger:** per-network
    - a metric for holding upstream networks accountable if they allow their users to congest downstream networks

BT

ECN
(recap)

code-point	standard designation
00	not-ECT
10	ECT(0)
01	ECT(1)
11	CE

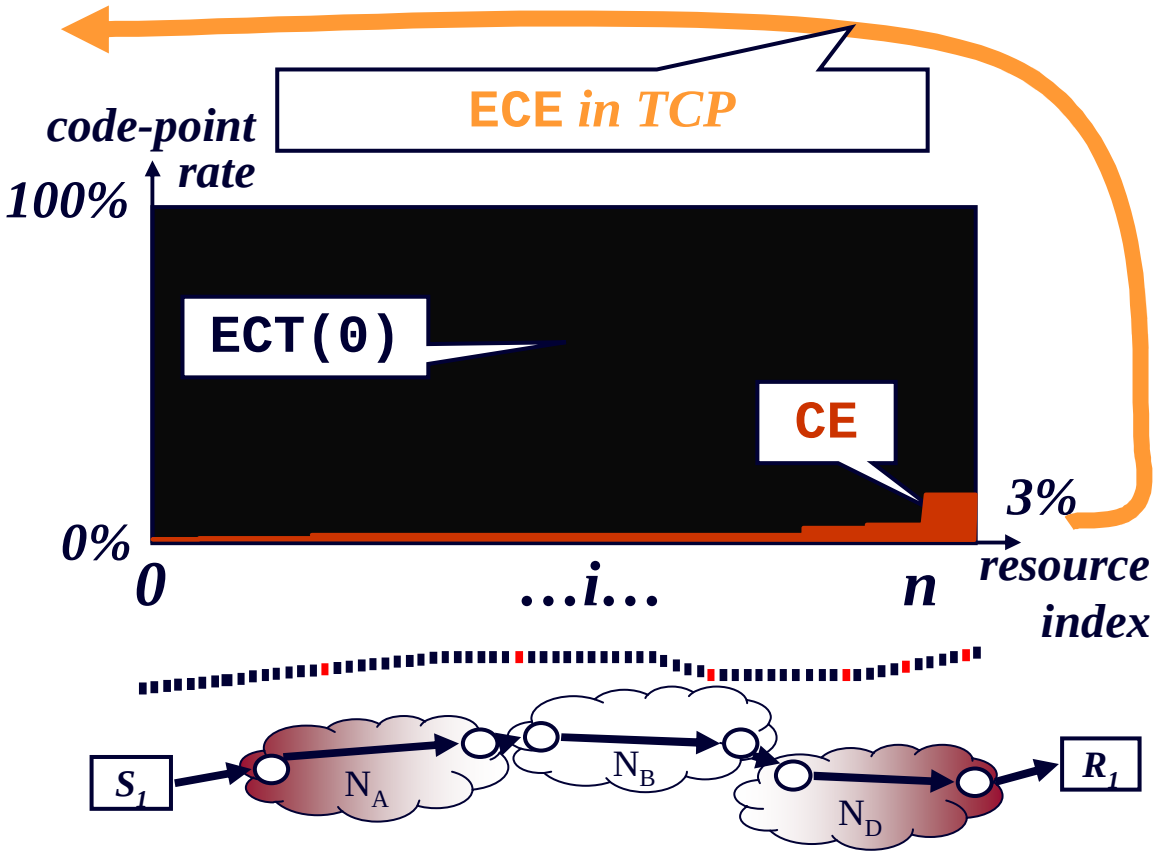

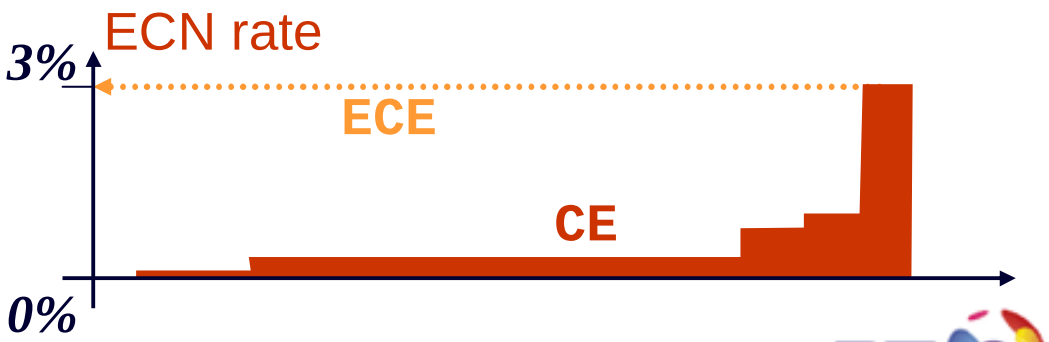

BT

re-ECN
(sketch)

code-point	standard designation
00	not-ECT
10	ECT(0)
01	ECT(1)
11	CE

- on every **Echo-CE** from TCP, sender sets **ECT(0)**, else sets ECT(1)
- at any point on path, diff betw rates of **ECT(0)** & **CE** is downstream congestion
- routers unchanged

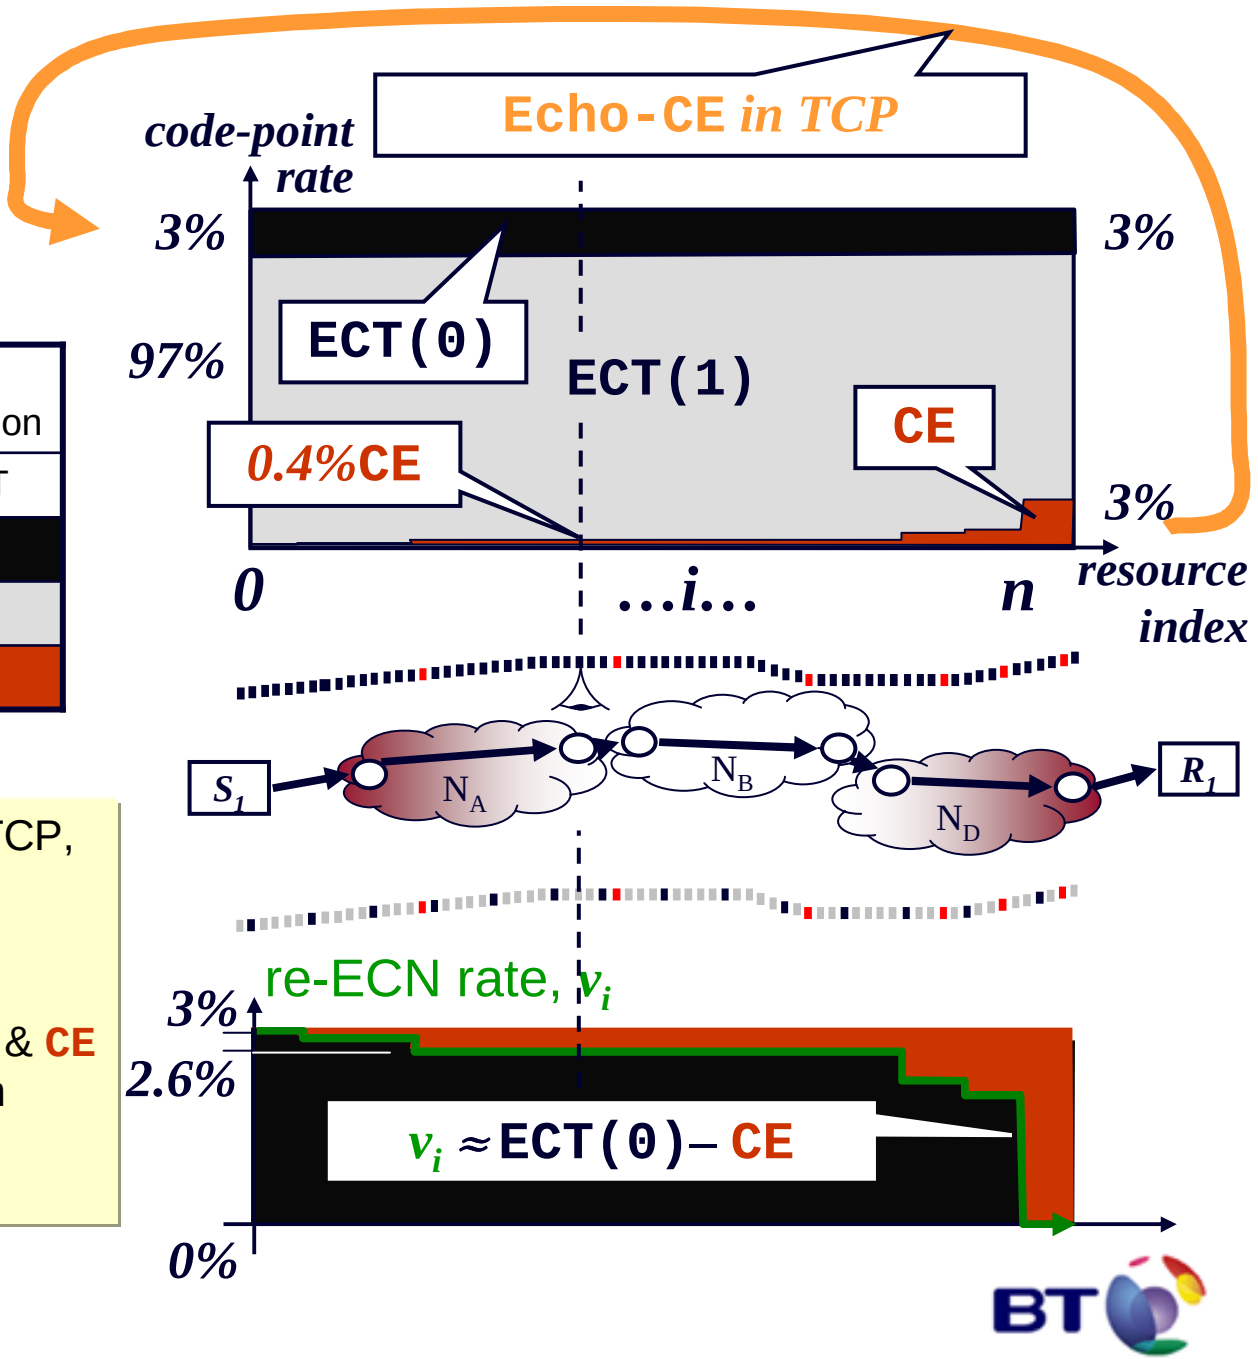

incentive framework
(user-network)

- packets carry view of downstream path congestion to each router
- so ingress can police rate response
  - using path congestion declared by sender
- won't snd or rcv just understate congestion?
- no – egress drops negative balance

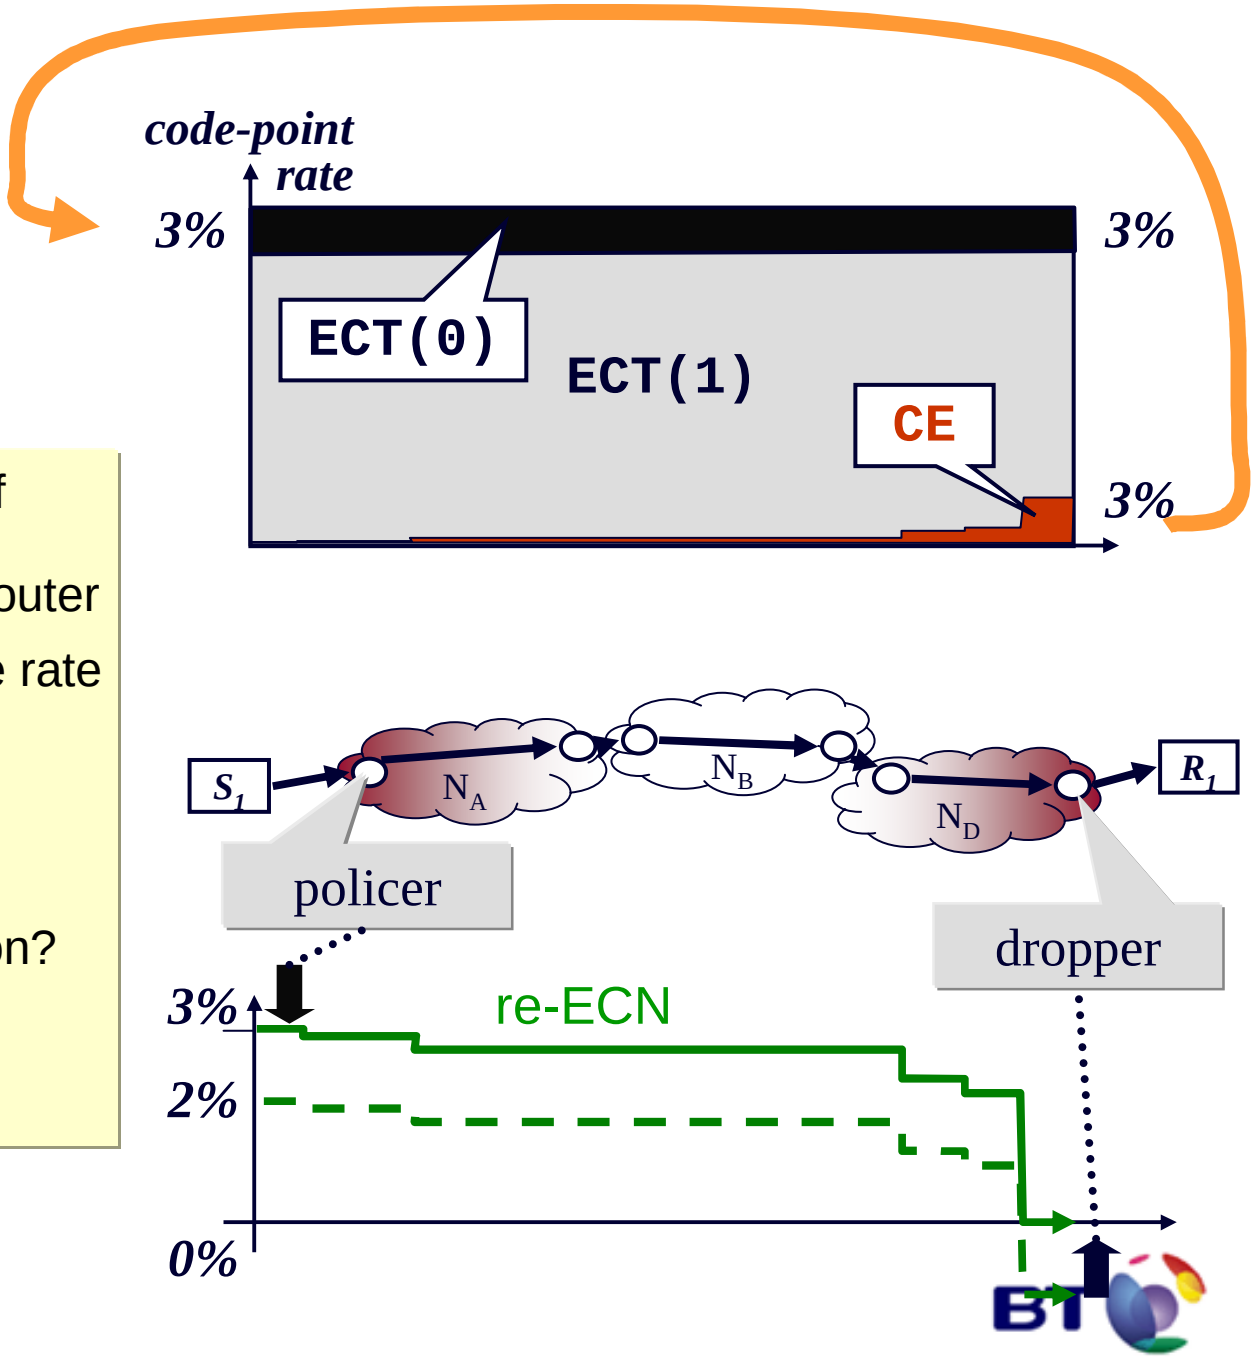

accountability for congestion

other applications

- congestion-history-based policer (congestion cap)
  - throttles causes of past heavy congestion (zombies, 24x7 p2p)
- DDoS mitigation
- QoS & DCCP profile flexibility
  - ingress can unilaterally allow different rate responses to congestion
- load sharing, traffic engineering
  - multipath routers can compare downstream congestion
- bulk metric for inter-domain SLAs or charges
  - bulk volume of **ECT(0)** less bulk volume of **CE**
  - upstream networks that do nothing about policing, DoS, zombies etc will break SLA or get charged more

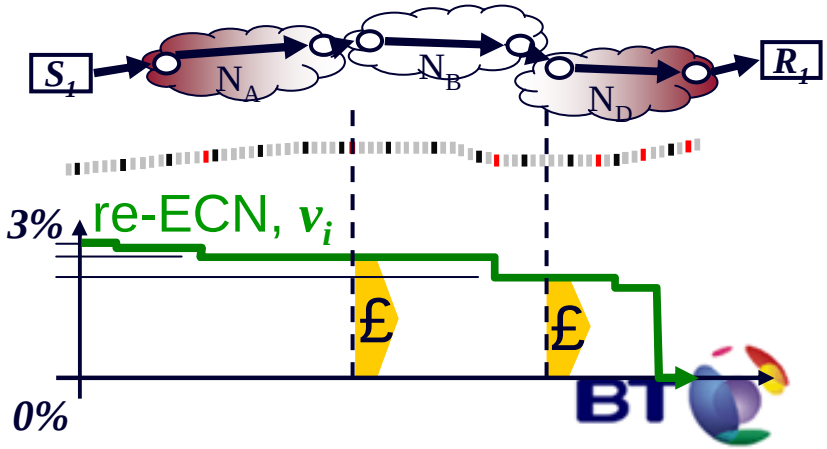

inter-domain accountability for congestion

- metric for inter-domain SLAs or charges
  - bulk volume of ECT(0)less bulk volume of CE
  - measure of downstream congestion allowed by upstream nets
  - volume charging tries to do this, but badly
  - aggregates and deaggregates precisely to responsible networks
  - upstream networks that do nothing about policing, DoS, zombies break SLA or get charged more

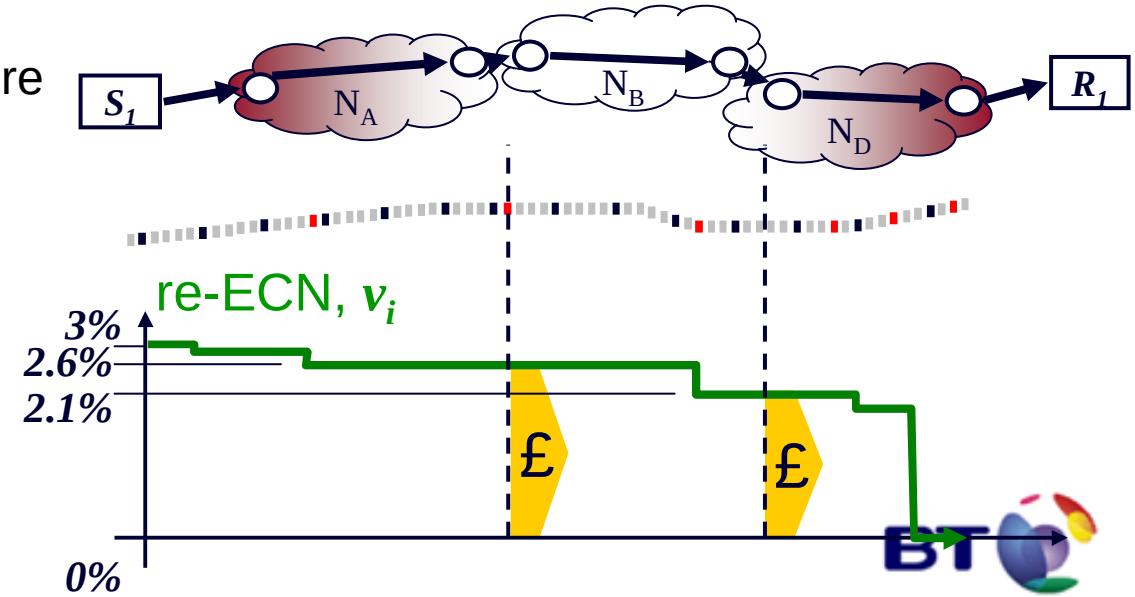

making stuff real

- tie to new product
  - the occasion when companies consider making changes
  - not just performance enhancement or cost reduction
- effort from inventors
  - not invented here has a flip side
  - hawking round every relevant forum – plan for long haul
  - creating a fashion

- unilateral action in the value chain
  - bilateral changes (e.g. vendor & operator) a second best
- bilateral between similar players (e.g. network operators)
  - bilateral between neighbours
  - overlays can turn remote networks into neighbours

BT

re-ECN incremental deployment

- only REQUIRED change is TCP sender behaviour
- precision only if receiver is re-ECN capable too
- optional compatibility mode for ‘legacy’ ECN rcvrs
  - inclined to leave it out (so few Legacy-ECN hosts out there)
- no change from ECN behaviour for
  - routers
  - tunnels
  - IPsec
  - middleboxes etc
- add egress droppers and ingress policers over time
  - policers not necessary in front of trusted senders

BT

re-ECN deployment transition

- if legacy firewalls block FE=1, sender falls back to FE=0
  - FE=0 on first packets anyway, so see connectivity before setting FE=1
  - if sender has to wrongly clear FE=0, makes dropper over-strict for all
- sender (and receiver): re-ECN transport (from legacy ECN)
  - ingress policer (deliberately) thinks legacy ECN is highly congested
    - 50% for nonce senders, 100% for legacy ECN
  - policers should initially be configured permissively
  - over time, making them stricter encourages upgrade from ECN to re-ECN

re-ECN deployment incentives - networks

- access network operators
  - revenue defence for their QoS products
  - can require competing streaming services over best efforts to buy the right to be unresponsive to congestion
- egress access operators: dropper
  - can hold upstream neighbour networks accountable for congestion they cause in egress access
  - without egress dropper, border congestion could be understated
- ingress access operators: policer
  - if downstream networks hold upstream accountable (above)
  - ingress will want to police its heavy & malicious users
  - ingress can choose to rate-limit Not-ECT
- backbone networks
  - unless hold upstream accountable will be held accountable by downstream

BT

re-ECN deployment incentives - vendors

- vendors of policing equipment
  - network operators invite to tender
- sender (and rcvr): re-ECN transport (from Not-ECT)
  - network operator pressure encourages OS vendor upgrades (sweetener below)
  - Not-ECT rate-limits (above) encourage user upgrades
- end device OS vendors
  - network operators hold levers (policers) to encourage customer product upgrades

everyone gains from adding accountability to TCP/IP
except the selfish and malicious

BT

making stuff real

- tie to new product
  - the occasion when companies consider making changes
  - not just performance enhancement or cost reduction
- effort from inventors
  - not invented here has a flip side
  - hawking round every relevant forum – plan for long haul
  - creating a fashion

- unilateral action in the value chain
  - bilateral changes (e.g. vendor & operator) a second best
- bilateral between similar players (e.g. network operators)
  - bilateral between neighbours
  - overlays can turn remote networks into neighbours

BT