

# re-feedback

opening a new chapter of Internet innovation?

**Bob Briscoe**  
Chief Researcher, BT  
Mar 2009



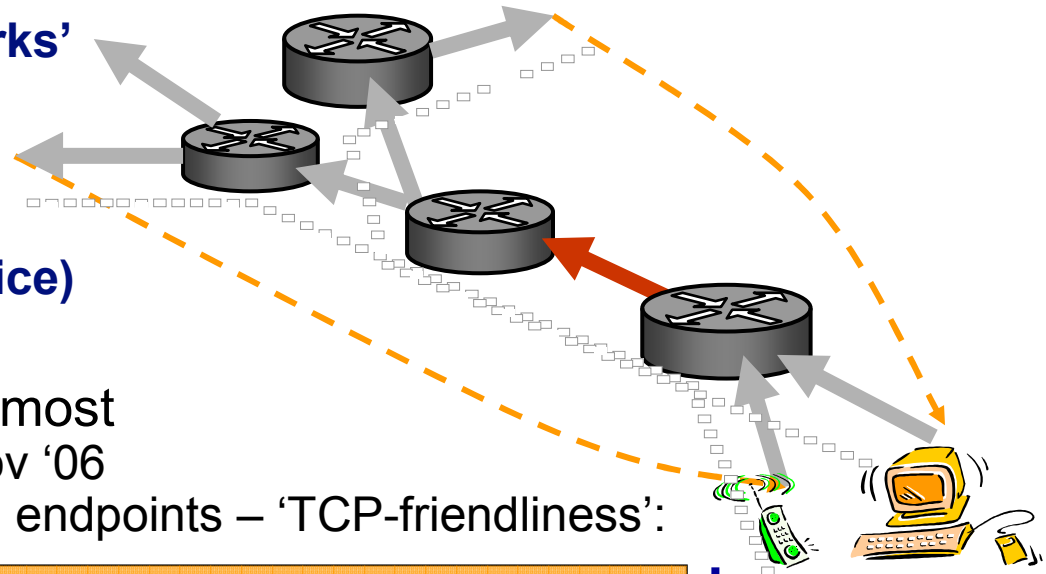
# how to share a cloud

- anyone can use any capacity anywhere on the Internet, as much as they like, without asking
  - fantastic ideal
  - but when freedoms collide, what share do you get?
- decades of misunderstanding
  - about accountability for usage costs of pooled bandwidth
- freedom without accountability?
  - operators would rather carve up resources than pool them
- proposed solution shares a cloud and caters for
  - self-interest & malice
    - of users and of providers
  - evolvability
    - of new application behaviours
    - of new business models
  - viability of supply chain
  - simplicity


# how classic Internet sharing 'works'

## endemic congestion & voluntary restraint

(increasingly overridden in practice)



- aka. those who take most, get most
  - technical consensus until Nov '06
  - voluntarily polite algorithm in endpoints – 'TCP-friendliness':



- a game of chicken – taking all and holding your ground pays
  - unresponsive flow<sub>3</sub> (VoIP, VoD, Joost 700kbps)




- or starting more 'TCP-friendly' flows than anyone else (Web: x2, p2p: x5-100)
- or for much much longer than anyone else (p2p file-sharing x200)
- poss. net effect of both (p2p: x1,000-20,000 higher traffic intensity)



# cost-shifting between services

- scenario
  - NGN sells unlimited Internet service at an average cost
  - NGN also a higher level service provider (TV, video phone)
  - competing with independent service providers (Skype, YouTube)
- who pays for capacity & QoS costs of higher value services?
  - NGN service layer must pay these costs internally
  - higher than average delivery cost of independent service
    - shifted to NGN's lighter users
    - because of how Internet sharing 'works'



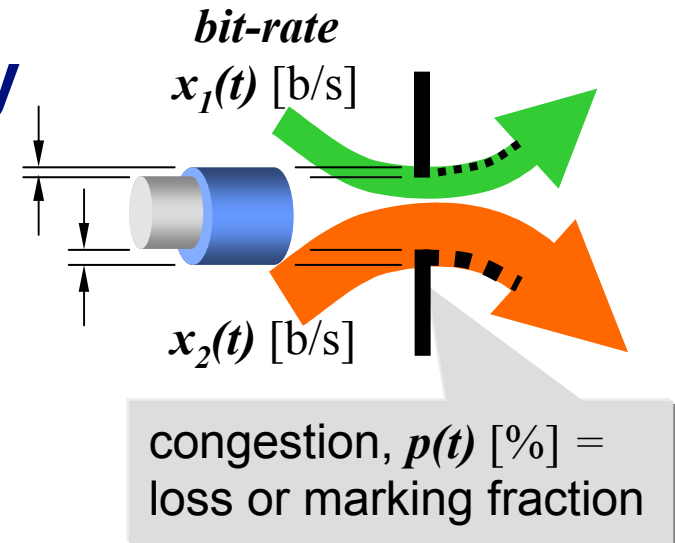
# blame our choices not the customers'

- network operator's knee-jerk reaction
  - throttle p2p or independent services
- but first, any CEO of a network business should be able to answer this question
  - Q. what is the cost of network usage?
  - A. volume? NO; bit-rate? NO
  - A. 'congestion-volume' (next slide)
- cannot be measured in current Internet protocol (IP) networks
- we haven't designed *our* contracts & *our* technology for machine-powered customers
  - attack customers' choices?
  - or attack the cost accountability deficiencies of IP?



# a new metric for accountability

## – a bandwidth trading unit



- congestion-volume
    - your volume weighted by congestion when it was sent
    - distinguishes friendly & hostile volume
  - intuition
    - some ISPs count volume during peak
    - like counting (100% x volume) during peak and (0% x volume) otherwise
    - congestion-volume =  $(p \times x_i \text{ over } time)$
  - how to measure
    - volume discarded from your traffic
    - more interesting: volume marked with explicit congestion notification (ECN)
- a resource accountability metric
    - of customers to ISPs (too much traffic)
    - and ISPs to customers (too little capacity)
  - 1. cost to other users of your traffic
  - 2. marginal cost of equipment upgrade
    - so it wouldn't have been congested
    - so traffic wouldn't have affected others
  - competitive market matches 1 & 2

*note: diagram is conceptual  
congestion volume & capital cost of equipment would be accumulated over time*



# if incoming congestion were visible congestion policing

## Acceptable Use Policy

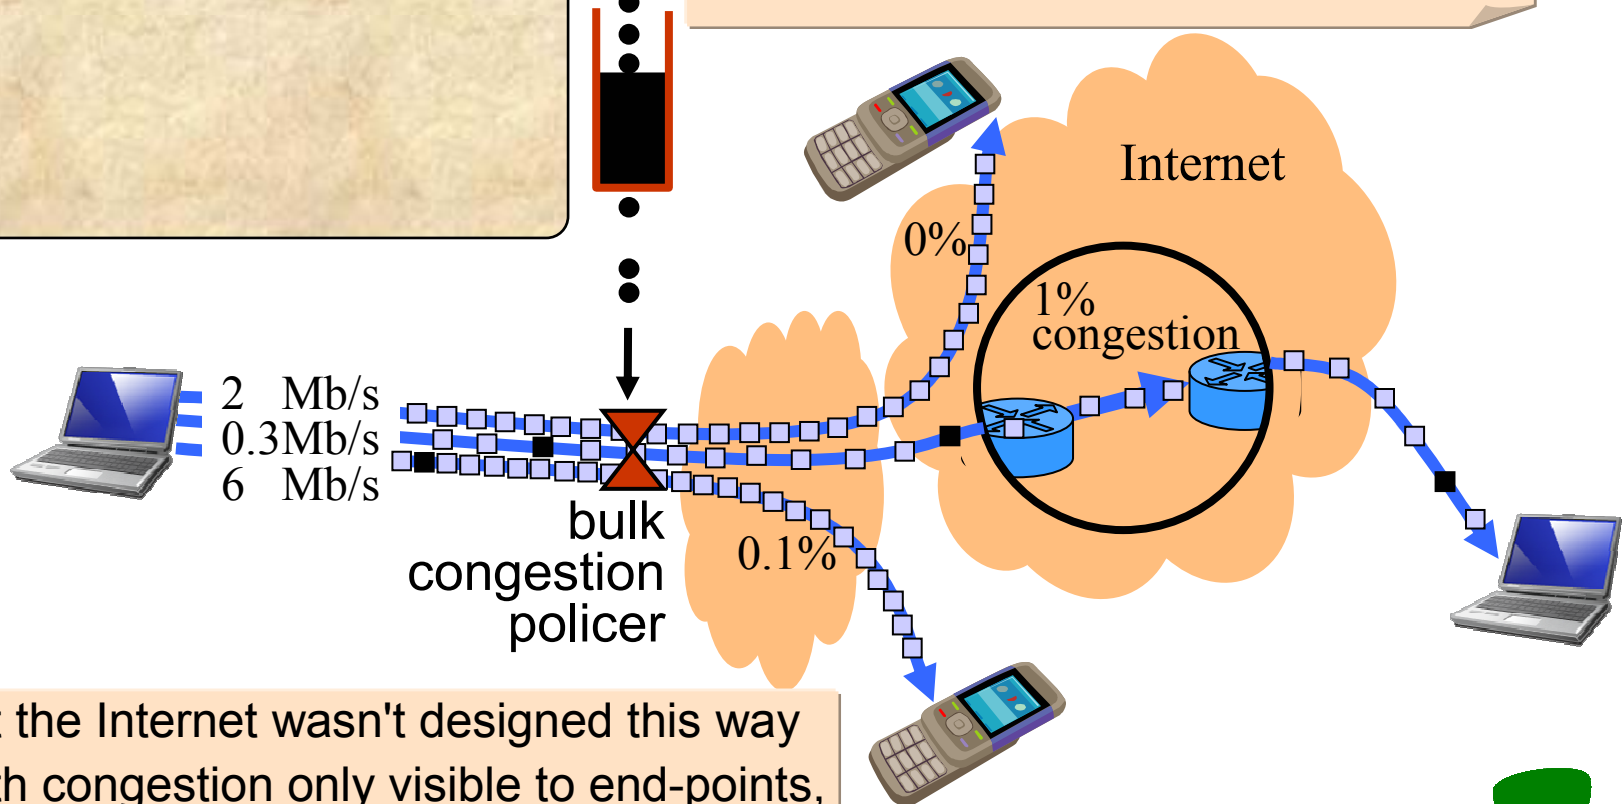
Your 'congestion-volume' allowance:

**1GB/month (= 3kb/s continuous)**

Only limits excess traffic above the Internet 'high-water-mark'

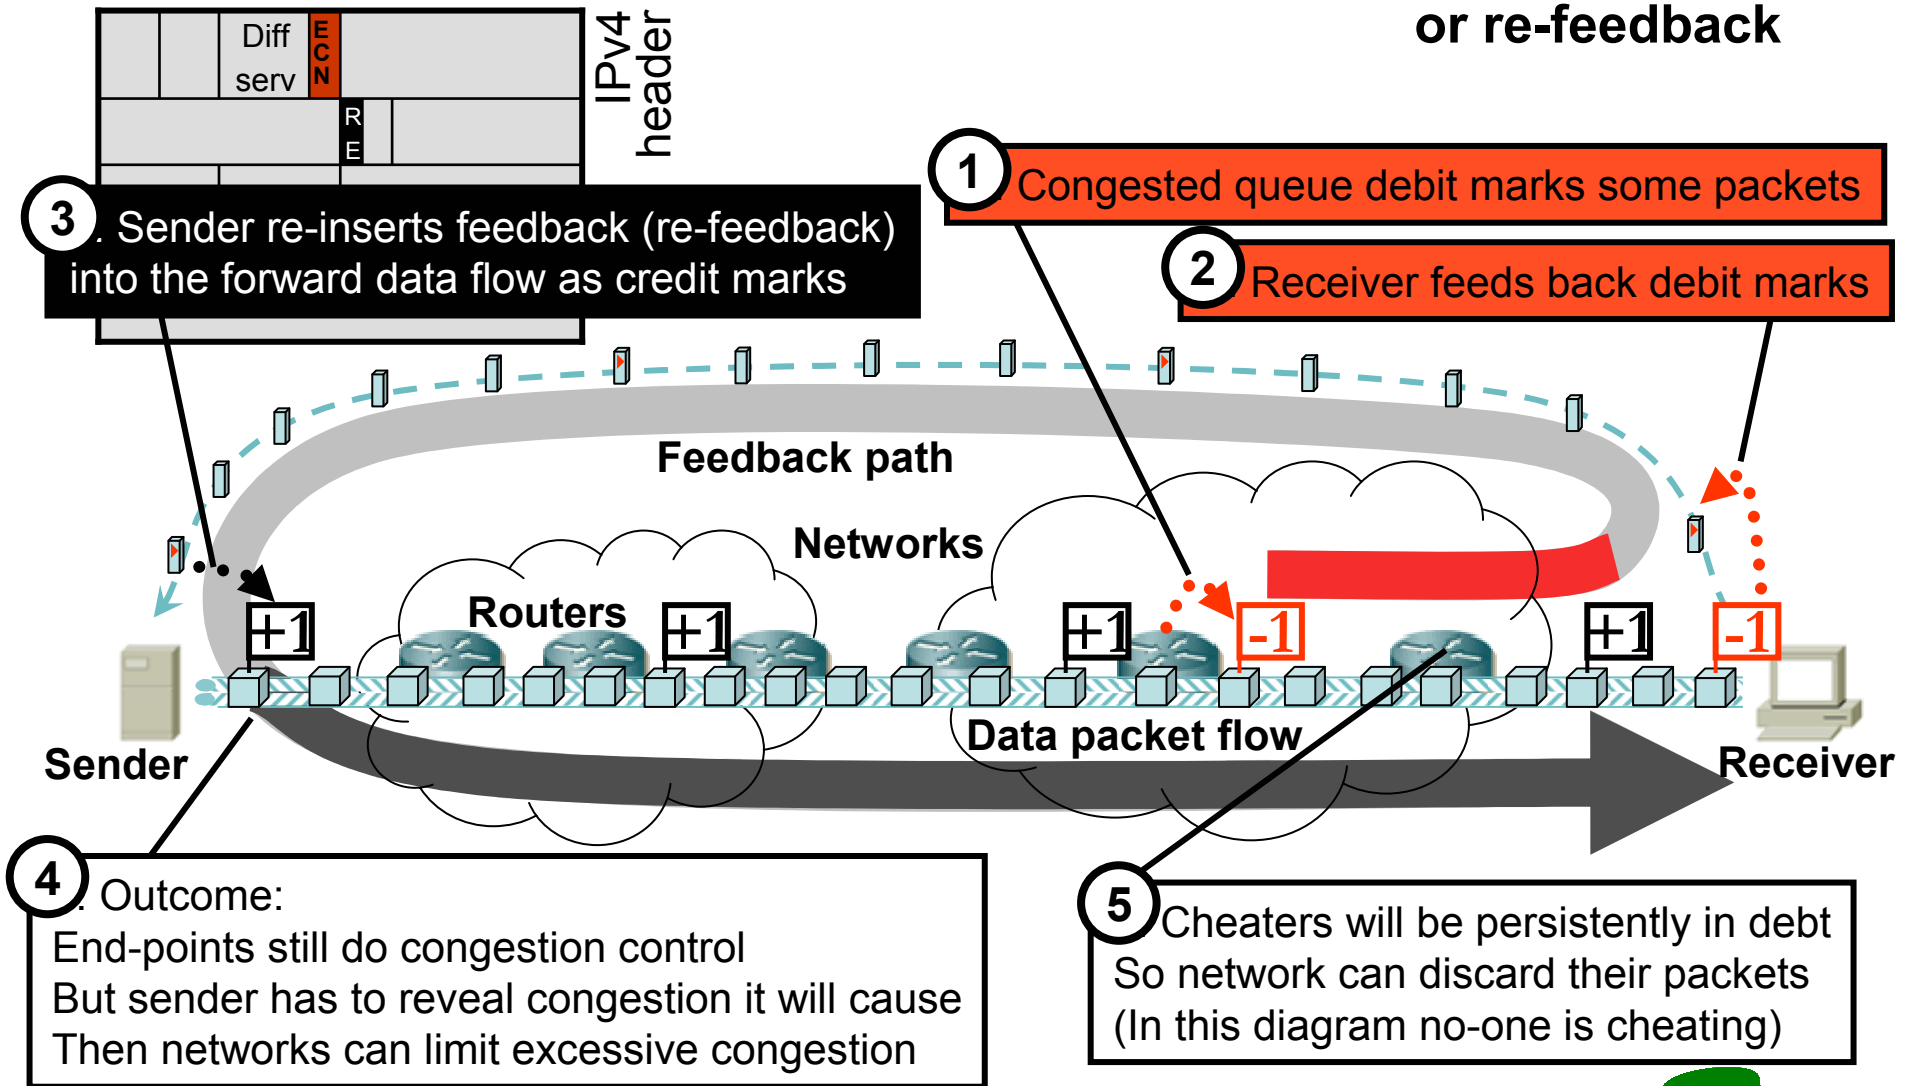
Under typical conditions this will allow you to transfer about **70GB per day**

- only throttles traffic when your contribution to congestion in the cloud exceeds your allowance



- but the Internet wasn't designed this way
- path congestion only visible to end-points, not to network

# re-ECN = standard ECN + re-inserted feedback or re-feedback




- No changes required to data forwarding




# interconnect aggregation

## simple internalisation of externalities

### 'routing money'



area =  
instantaneous  
downstream  
congestion  
volume

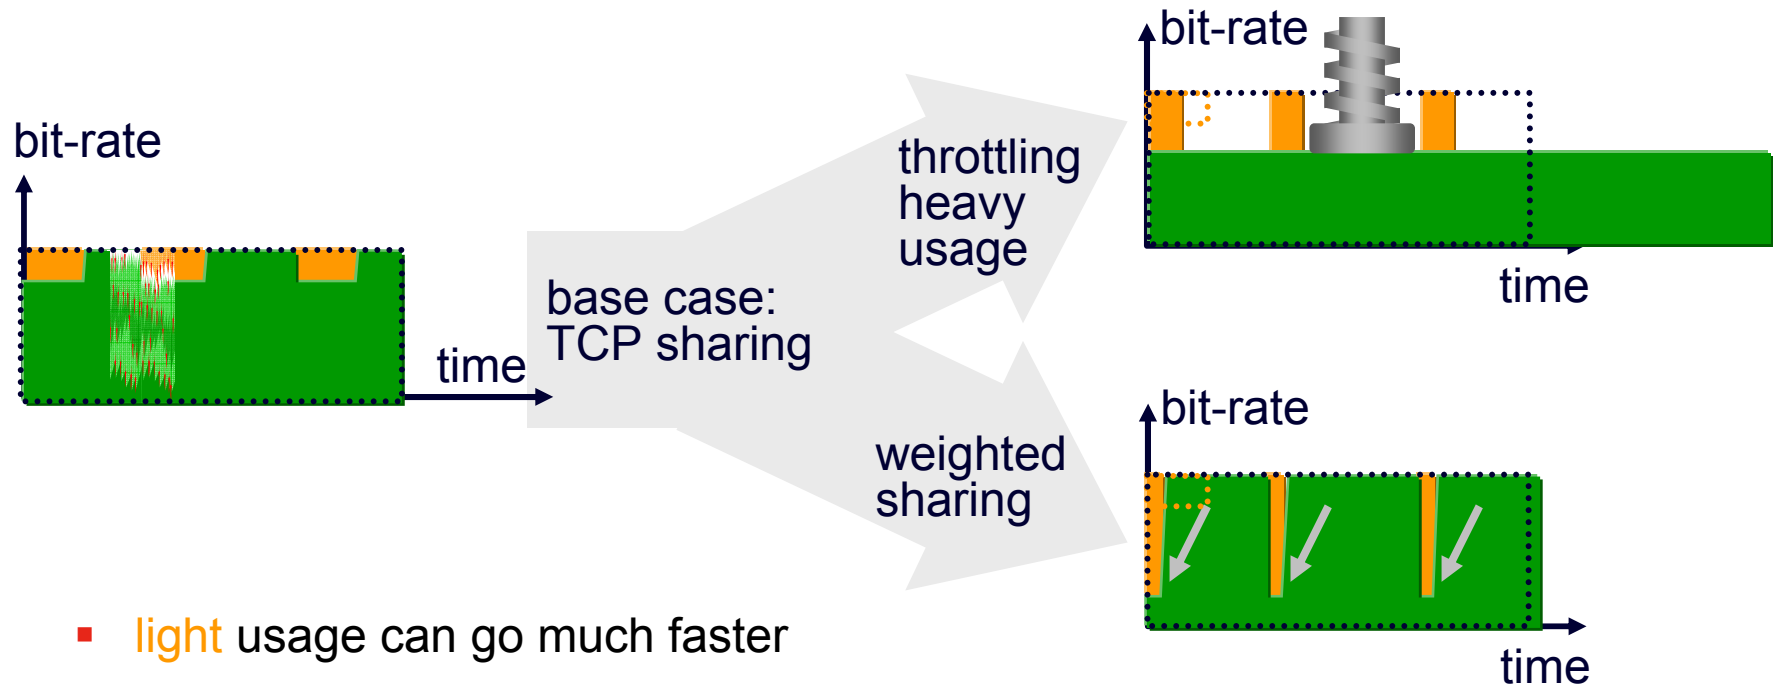


0|0|2|7|6|0|5

just two counters at border  
meter monthly bulk volume  
of marked packets  
= aggregate downstream  
congestion-volume in flows  
*without measuring flows*



# there are better solutions than fighting



- **light** usage can go much faster
- hardly affecting completion times of **heavy** usage
- only requires bulk congestion policer incentives
  - evolution path to simpler quality of service, with multipath, multi-site, roaming,...
  - denial of service mitigation


NOTE: weighted sharing doesn't nec. imply differentiated services

- can be weighted aggressiveness of end-point rate control
- 'Design for Tussle' advises allow both



# openness – a tactic not a strategy

- once true usage costs visible at ingress control point
- retailers can choose
  - how tightly to control their customers' choices vs cost
  - their market position between open and closed
- changing your mind
  - involves changing a policy
  - not new technology
- truly converged architecture
  - for vendors & wholesalers
  - but no need for retail monoculture
- designed for tussle



# a new chapter of innovation

- applications & services
  - opens whole new space
- transport layer on end-points
  - usage costs currently visible here
- internetwork layer
  - once usage costs revealed here
  - ISPs won't need deep packet inspection for cost control
- link layer
  - once network layer can limit congestion
  - can remove bit-rate limits in shared access networks: passive optical networks, cable, wireless, cellular...



# re-feedback

opening a new chapter of Internet innovation?

clarification questions?  
spare slides


triple



# problems using congestion in contracts

	1. loss	2. ECN	3. re-ECN
can't justify selling an impairment	☹	☺	☺
absence of packets is not a contractible metric	☹	☺	☺
congestion not visible to upstream network nodes	☹	☹	☺
congestion is outside a customer's control	☹	☹	☺
customers don't like variable charges	☹	☹	☺
congestion is not an intuitive contractual metric	☹	☹	☹

- 1. loss:** used to signal congestion since the Internet's inception
  - computers detect congestion by detecting gaps in the sequence of packets
  - computers can hide these gaps from the network with encryption
- 2. explicit congestion notification (ECN):** standardised into TCP/IP in 2001
  - approaching congestion, a link marks an increasing fraction of packets
  - implemented in Windows Vista (but off by default) and Linux, and IP routers (often off by default)



- 3. re-inserted ECN (re-ECN):** standards proposal since 2005 (later slides)
  - packet delivery conditional on sender declaring expected congestion
  - uses ECN equipment in the network unchanged



# main steps to deploy re-feedback / re-ECN

- protocol
  - assign one bit in Internet protocol to re-ECN
- network (incremental)
  - turn on explicit congestion notification in routers (already available)
  - deploy simple policing functions at customer interfaces around participating networks
  - passive metering functions at inter-domain borders
- terminal devices (incremental)
  - (minor) addition to TCP/IP stack of sending device
  - or sender proxy in network
- contracts (incremental)
  - congestion-volume allowance in customer contracts
  - use congestion-volume in interconnection contracts

