

Internet cost transparency

mending value chain incentives

Bob Briscoe
Chief Researcher, BT
Sep 2009

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



capacity sharing


- raison d'être of the Internet
 - not just core & regional backhaul
 - shared access: wireless, cable, optical
- anyone can take any share of any link in the Internet
 - fantastic ideal
 - but when freedoms collide, what share do you get?

how to share Internet capacity?

- the invisible hand of the market, whether competitive or regulated
 - favours ISPs that share capacity in their customers' best interests

- since 1988 misplaced belief in TCP alone as the sharing standard


- ISP's homespun alternatives have silently overridden TCP
 - ad hoc application-specific blocks and permits
 - deep packet inspection
 - nailed up capacity
 - ...




how Internet sharing 'works'

TCP-fairness


- voluntarily polite algorithms in endpoints
 - pushes until congested
 - equalises rates of data flows



a game of chicken – taking all and holding your ground pays



or start more 'TCP-fair' flows than anyone else (**Web: x2, p2p: x5-100**)




or for much more data than others (**video streaming or p2p file-sharing x200**)

- net effect of both (p2p: x1,000-20,000 higher traffic intensity)

ISP's homespun alternatives have silently overridden TCP


1. equal bottleneck flow rates (TCP)?
2. access rate shared between active users, but weighted by fee (weighed fair queuing, WFQ)?
3. volume caps tiered by fee?
4. heaviest applications of heaviest users throttled at peak times by deep packet inspection (DPI)?




no current solution

harnesses end-system flexibility


1. TCP




2. (weighted) fair queuing



3. volume caps




4. deep packet inspection (DPI)



simpler & better...



weighted
TCP
sharing



- **light** usage can go much faster
- hardly affects completion time of **heavy** usage
 - doesn't have to shift into night
- BitTorrent & Microsoft have protocols to do this

but... punished by #2, #3 & #4

NOTE: weighted sharing doesn't imply differentiated network service

- just weighted aggressiveness of end-system's rate response to congestion



closing off the future

- becoming impossible to deploy a new use of the Internet
 - must negotiate arbitrary blocks and throttles en route
- two confusable motives
 - fairer cost sharing
 - competitive advantage to own services
- how to deconfuse? how to encourage fairer cost sharing?
 - make cost of usage transparent
- fixing Internet technology should avoid need for legislation

the missing link




Q. what is the marginal cost of a customer's usage?

A. each customer's contribution to congestion
congestion-volume

- unforgivable for a network business not to understand its primary marginal cost


isn't volume a good enough cost metric?



congestion is not evil

congestion *signals* are healthy

- no congestion across whole path is evil
 - for data transfer to complete ASAP, must fill bottlenecks







the trick

signal congestion *just before* impairment

- explicit congestion notification (ECN)
 - 2001 update to IP: as a queue builds mark more packets
- then tiny queuing delay and tiny loss for all traffic

measuring marginal cost

- user's contribution to congestion
= bytes marked
- can transfer v high volume
 - but keep congestion-volume v low
 - similar trick for video streaming




congestion-volume metric dual demand & supply role

- a resource accountability metric
 1. of customers to ISPs (too much traffic)
 2. and ISPs to customers (too little capacity)

- 1. cost to other users of my traffic
- 2. the marginal cost of upgrading equipment
 - so it wouldn't have been congested

- competitive market matches 1 & 2


- customer tells ISP which demand is worthy of capacity investment



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

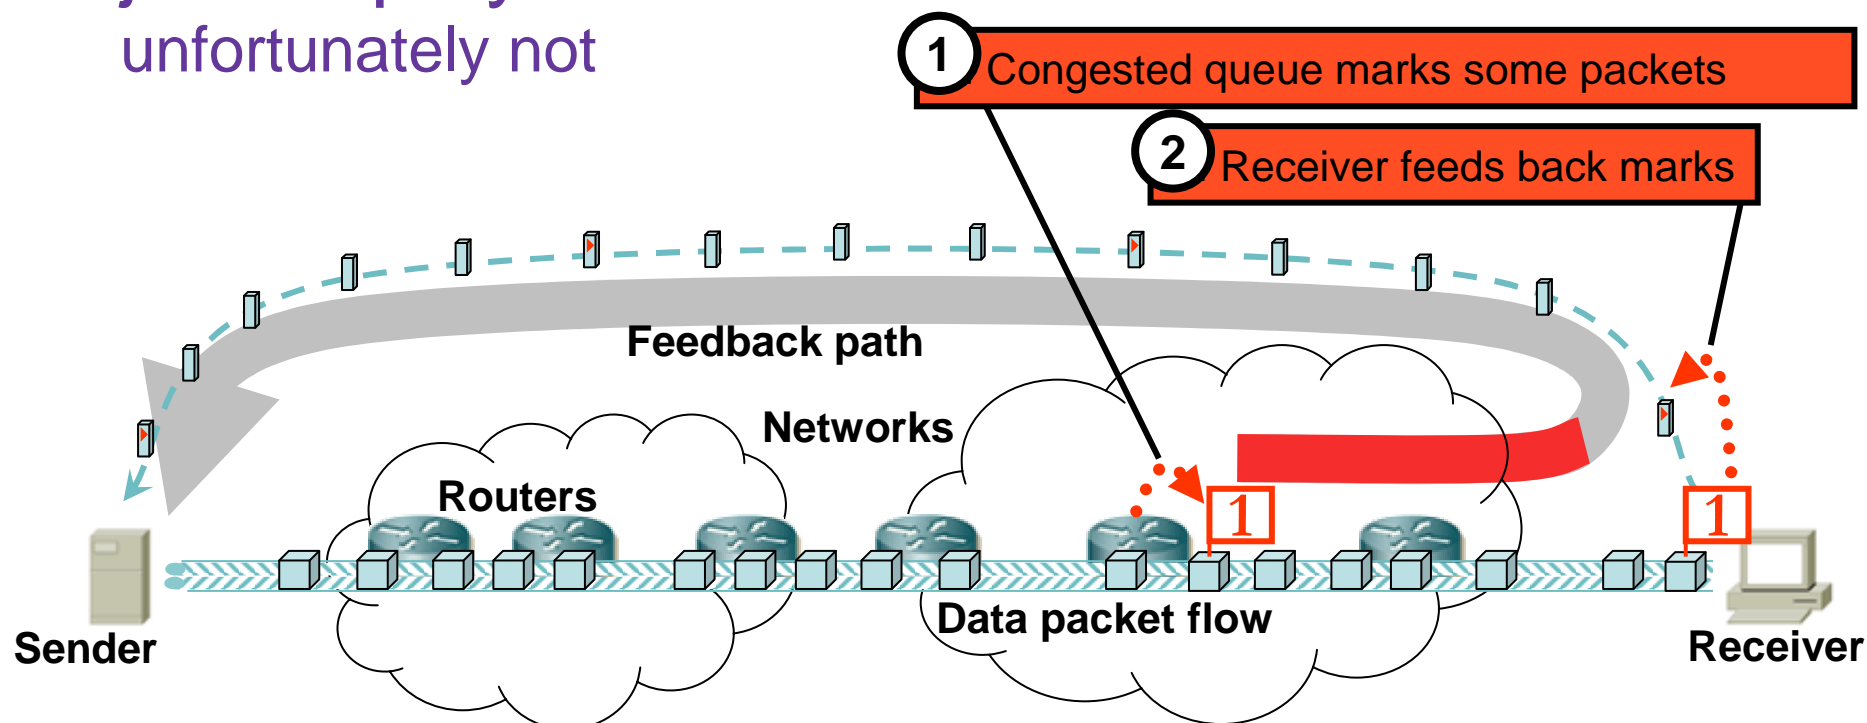
root cause

- by Internet design
 - end-systems manage congestion
- fine, but ISPs need to see it too
 - “cost transparency”
- ISPs cannot see primary business metric
 - packet loss can certainly be measured locally
 - but not a robust contractual metric – an absence & an impairment
- lacking visibility of congestion, ISPs:
 - punish nice and nasty volume equally
 - block light usage from going fast, even momentarily
 - require high cost apps (VoD, etc) to seek permission



just deploy ECN and we're done?

unfortunately not




- can only count ECN received, not sent
 - sender controls how much congestion receiver receives
 - consequence of Internet's one-way datagram model
- incentives would all be backwards
 - for receivers & for receiving networks

summary so far

the problem

- everyone thought fairness goal was equal flow rates
 - didn't take account of range of users' data activity over time
- ISPs trying to pull system to a different allocation
 - lacking visibility of the marginal costs
 - resorting to means confusable with non-neutrality



Internet cost transparency

proposed
solution



proposed solution

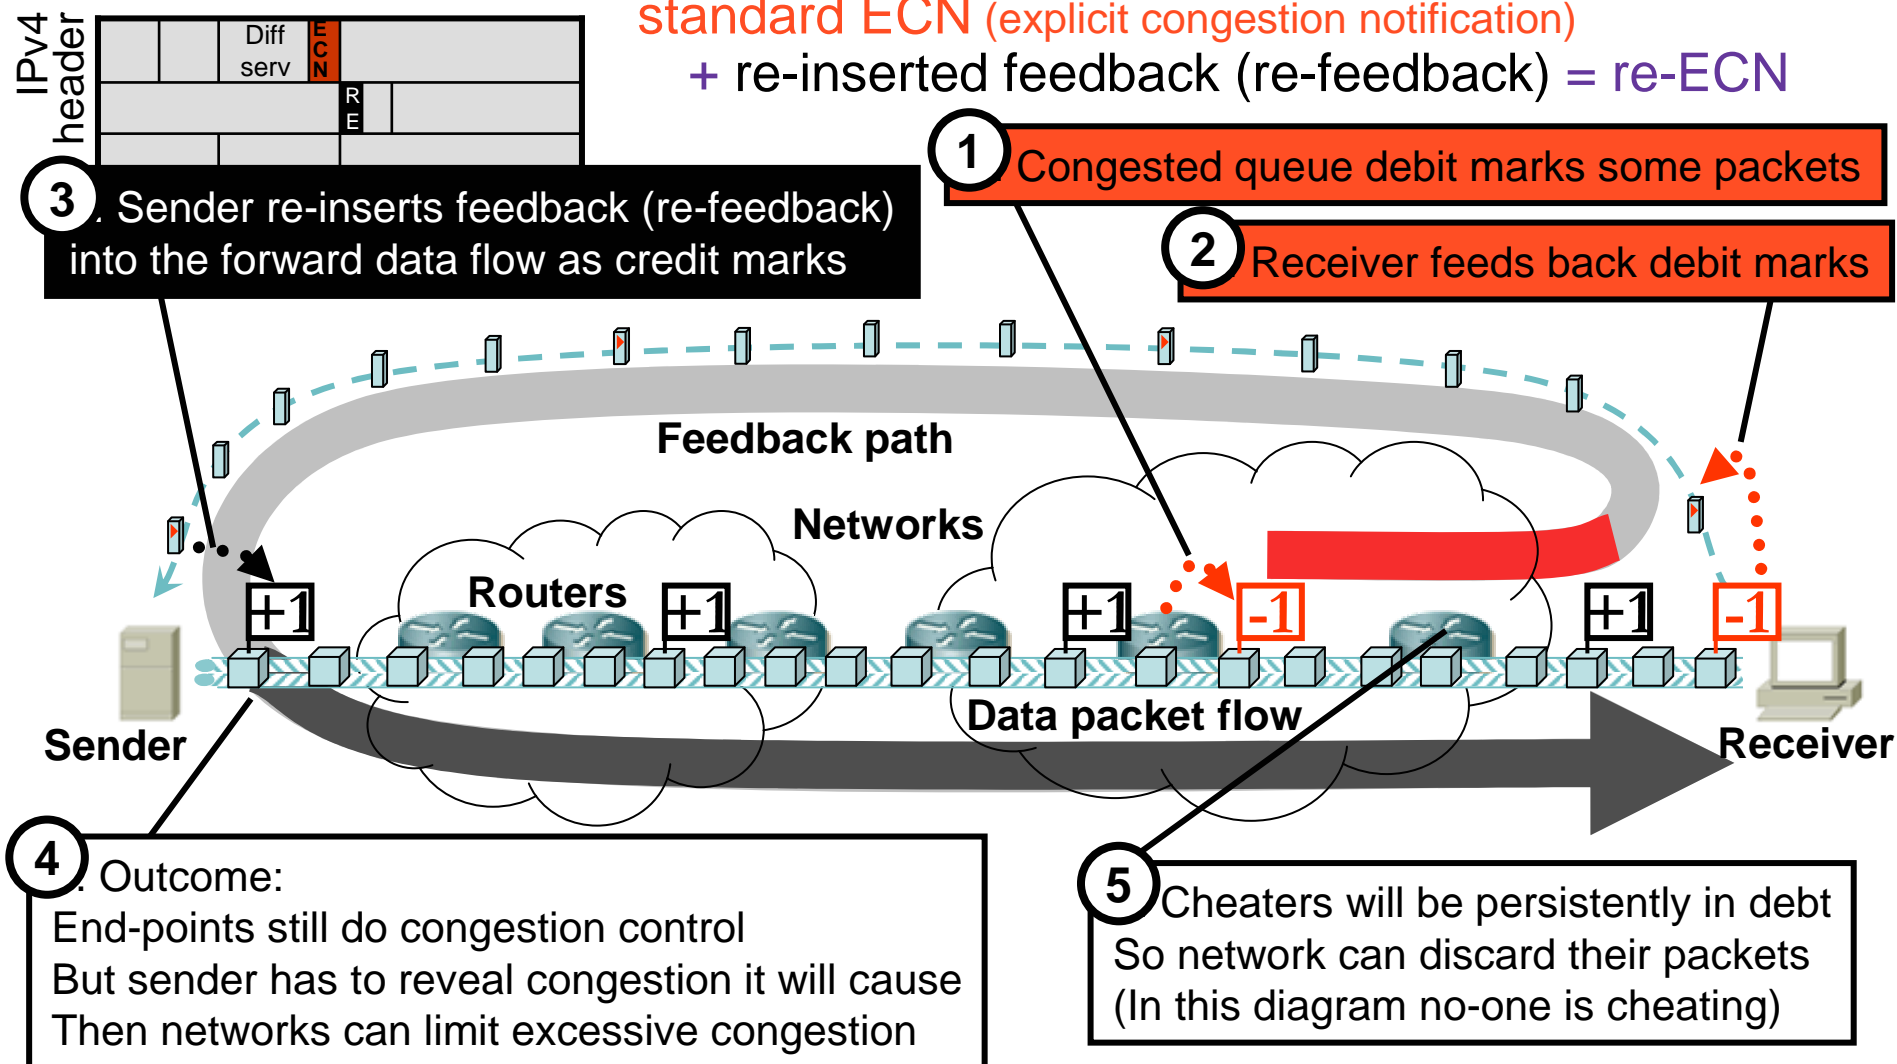
- mechanism & incentive
 - for sender to reveal congestion to network
 - so ISP can count contribution to congestion as easily as volume
- easy to build accountability models on top
 - accountability of customer to ISP
 - ISP to customer
 - ISP to ISP
- should greatly simplify operational support systems

one bit opens up the future



standard ECN (explicit congestion notification)

+ re-inserted feedback (re-feedback) = re-ECN



no changes required to IP data forwarding

status – IETF

glossary

IETF Internet Engineering Task Force

IESG Internet Engineering Steering Group

IAB Internet Architecture Board

IRTF Internet Research Task Force



- since 2006 IETF support for TCP capacity sharing has collapsed to zero
 - thought leaders agree TCP dynamics correct, but sharing goal wrong
 - many support our new direction – not universally – yet!
 - rewrite of IETF capacity sharing architecture in process
 - IETF delegated process to IRTF design team
- early Sep'09
 - proposed IETF working group: “congestion exposure” (experimental)
 - >40 offers of significant help in last fortnight
 - Microsoft, Nokia, Cisco, Huawei, Alcatel-Lucent, NEC, Ericsson, NSN, Sandvine, Comcast, Verizon, ...
 - 2 days ago: IESG / IAB allowed agenda time, Hiroshima Nov'09
 - non-binding vote on working group formation
- not a decision to change to IP – defer until support is much wider

I E T F[®]

example#1: retail flat fee congestion policing



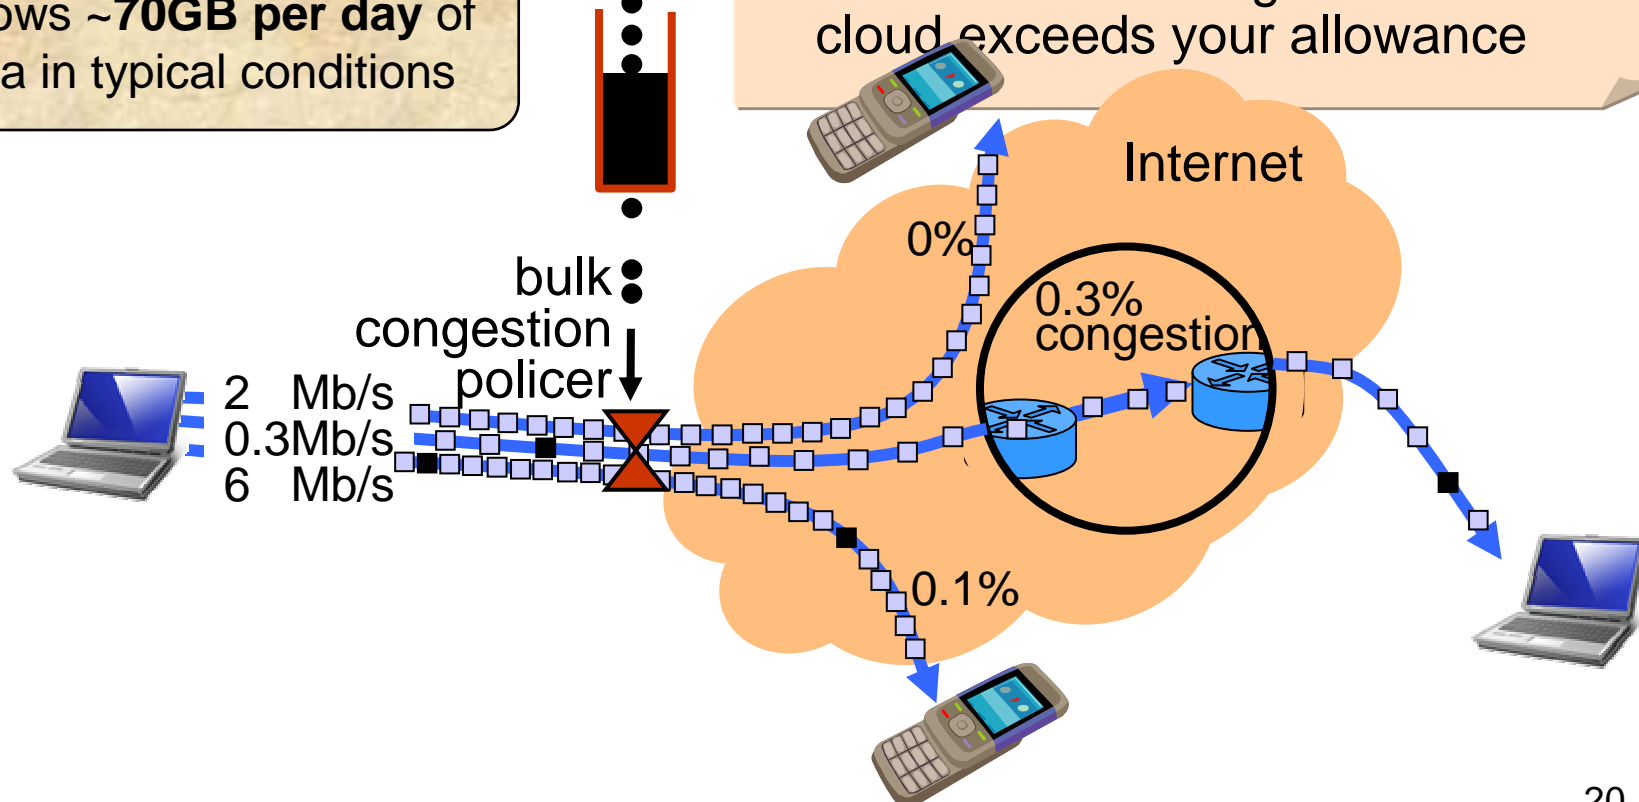
Acceptable Use Policy

'congestion-volume'
allowance: 1GB/month


@ €15/month

Allows ~70GB per day of
data in typical conditions

- simple invisible QoS mechanism
 - apps that need more, just go faster
- only throttles traffic when your contribution to congestion in the cloud exceeds your allowance



example#2: bulk cost in border SLAs 'routing money'





bringing information to the control point

- flat fee policer is just one example...
- huge space for business & technical innovation at the control point

- cost based, value-cost based
- bulk, per flow, per session
- call admission control
- policing, charging
- tiers, continuous
- wholesale, retail

- truly converged architecture
 - can apply different industry cultures
 - through policies at the control point
 - not embedded in each technology



main steps to deploy re-feedback / re-ECN



summary

rather than control sharing in the access links,
pass congestion info & control upwards

- network
 - turn on explicit congestion notification in data forwarding
 - already standardised in IP & MPLS
 - standards required for meshed network technologies at layer 2 (ECN in IP sufficient for point to point links)
 - deploy simple active policing functions at customer interfaces around participating networks
 - passive metering functions at inter-domain borders
- terminal devices
 - (minor) addition to TCP/IP stack of sending device
 - or sender proxy in network
- then new phase of Internet evolution can start
 - customer contracts & interconnect contracts
 - endpoint applications and transports
- requires update to the IP standard (v4 & v6)
 - started process in Autumn 2005
 - using last available bit in IPv4 header or IPv6 extension header

summary

- the invisible hand of the market, whether competitive or regulated
 - favours ISPs that share capacity in their customers' best interests
- cost (congestion) transparency
 - customers reveal costs to providers
- aligns incentives
 1. primary Internet capacity sharing mechanism (weighted TCP)
 2. ISP policing mechanisms
- encourages diversity in both
- ensures whole value chain accounts for infrastructure costs
- a technology can't enforce neutrality
 1. can at least provide the means to run a viable neutral business in a commodity market
 2. for value-based business: reveals currently unknown costs
- joins up broken Internet value chain



content industry, CDNs, network wholesalers & retailers, Internet companies, end-customers, business, residential

more info...

- The whole story in 7 pages
 - Bob Briscoe, "[Internet Fairer is Faster](#)", BT White Paper (Jun 2009)
 - ...this formed the basis of:
 - Bob Briscoe, "[A Fairer, Faster Internet Protocol](#)", IEEE Spectrum (Dec 2008)
- Inevitability of policing
 - [CFP06] The Broadband Incentives Problem, Broadband Working Group, MIT, BT, Cisco, Comcast, Deutsche Telekom / T-Mobile, France Telecom, Intel, Motorola, Nokia, Nortel (May '05 & follow-up Jul '06) <cfp.mit.edu>
- Slaying myths about fair sharing of capacity
 - [Briscoe07] Bob Briscoe, "[Flow Rate Fairness: Dismantling a Religion](#)" ACM Computer Communications Review 37(2) 63-74 (Apr 2007)
- How wrong Internet capacity sharing is and why it's causing an arms race
 - Bob Briscoe et al, "[Problem Statement: Transport Protocols Don't Have To Do Fairness](#)", IETF Internet Draft (Jul 2008)
- Understanding why QoS interconnect is better understood as a congestion issue
 - Bob Briscoe and Steve Rudkin "[Commercial Models for IP Quality of Service Interconnect](#)" BT Technology Journal 23 (2) pp. 171--195 (April, 2005)
- Equitable quality video streaming
 - [Crabtree09] B. Crabtree, M. Nilsson, P. Mulroy and S. Appleby "Equitable quality video streaming" Computer Communications and Networking Conference, Las Vegas, (January 2009)

available from the re-ECN & re-feedback project page:

<http://bobbriscoe.net/projects/refb/>

bob.briscoe@bt.com

Internet cost transparency

Q&A...

& spare slides...



partial deployment of re-feedback / re-ECN

- network equipment
 - both policing & forwarding: each network that wants to see congestion can deploy independently of others
 - not all forwarding equipment can do ECN today
fine if it drops instead, esp if not frequently congested
- sender
 - distinction between re-ECN & non-re-ECN packets
 - sender can choose which it sends
 - if network is policing based on re-ECN info
it's likely to rate-limit non-re-ECN packets
- receiver
 - works OK with Vista receiver now
 - upgrade to receiver to work precisely

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 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 (off by default)

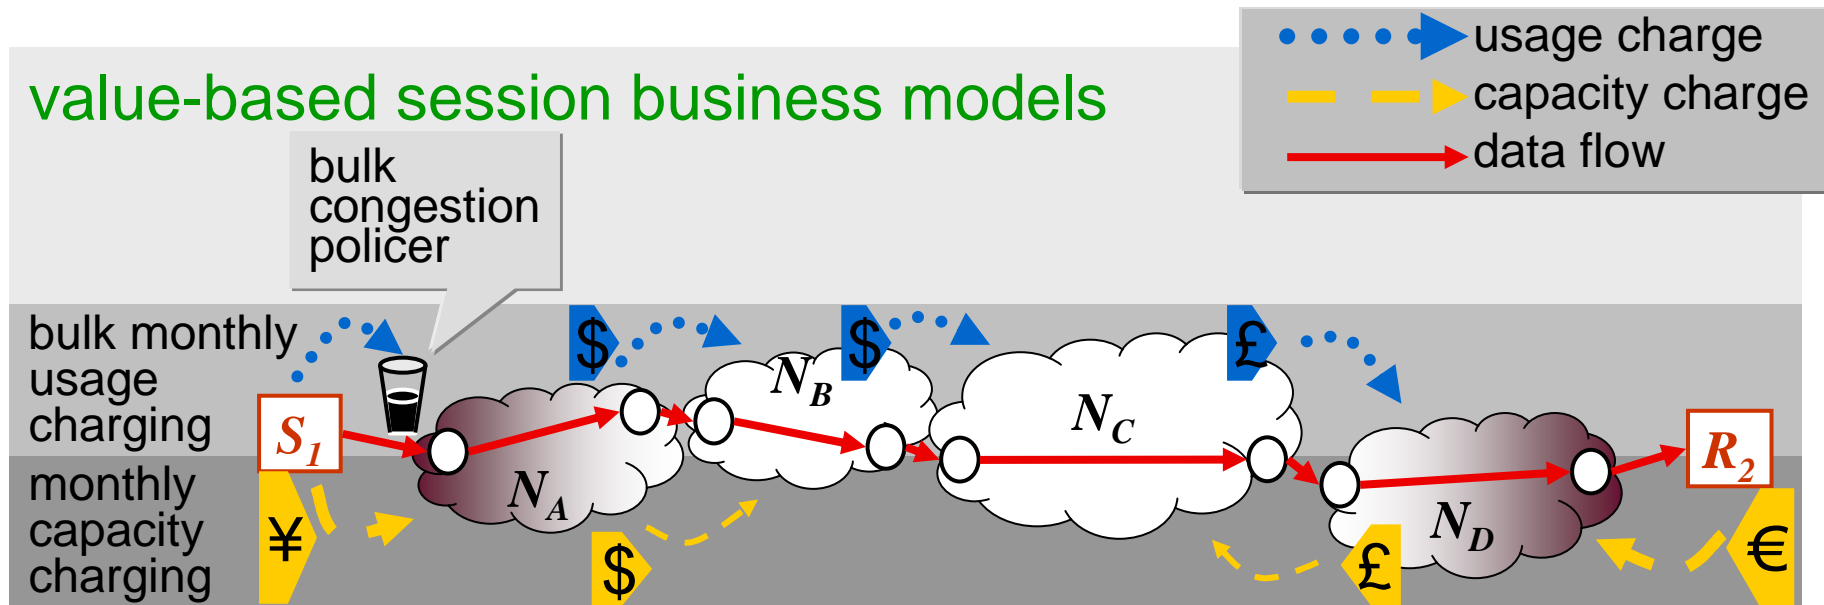


3. **re-inserted ECN (re-ECN)**: standards proposal since 2005


- packet delivery conditional on sender declaring expected congestion
- uses ECN equipment in the network unchanged

example sustainable business model for basic data transport

value-based session business models



can then be built (and destroyed) over this



a new chapter of innovation

- applications & services
- 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
 - can remove bit-rate limits in shared access: passive optical, cable, wireless, cellular...

novel service & app behaviours

battery optimisation

smooth quality video
>2x more videos

server DDoS protection

resilience
using multi-paths

QoS mechanism
simple – just go faster

hi-speed transfers

low latency
always

QoS interconnect
trivial

background transfers
incentivised

commercially viable interface to Internet layer

traffic engin'g
intra & inter

congestion
policing

access unbundling
at IP layer!

network DDoS
natural protection

shared medium access
delegate upwards

simpler access
technologies
potential