

Capacity Sharing: IETF Activities

NOKIA

Lars Eggert
Nokia Research Center

Bandwidth Bandwagon: An ISOC Briefing Panel
Hiroshima, Japan
November 11, 2009

The Internet is all about capacity sharing

connection-less

no isolation between flows

best-effort

no strict delivery guarantees

end-to-end

smart edges, dumb core

result: dramatic scaling potential
that enabled entire new industries

hard to imagine how another
approach could have resulted in
similar bang-for-the-buck



And sharing means caring

packets belonging to different flows share the path – and the behavior of each affects the others

- congestion


- queuing delay

apps & protos need “social” behavior

- otherwise the Internet stops

- being a useful shared resource

the IETF provides a toolbox of mechanisms that allow apps to share capacity intelligently



Architectural principles

remember: smart edges, dumb core


means that – in general – the responsibility is **split** between the apps & the network

network provides app- & service-agnostic information about path conditions in a timely manner

incl. loss = “am really overloaded”

apps (or the transport protocols they use) act on this information

how to act can be app-specific



So what is in the IETF toolbox?

Congestion control: TCP friendly

TCP + TCP friendly congestion control
hosts: determine transmission rate
according to path conditions based
on observed RTT and loss

optimizations/extensions:

Explicit Congestion Notification (ECN)

net: mark before drop

hosts: react to marks as if loss

Active Queue Management (e.g., RED)

net: progressive mark/drop



Low extra delay background transport

goal: transmit bulk data without substantially affecting the delay seen by other users and apps

approach: congestion control to:

- saturate the bottleneck = bulk data

- maintain a low extra delay

- yield to standard TCP = background

combine with less-than-best-effort

DiffServ + ISP pricing to give additional incentives for deployment

Fri 9:00 LEDBAT WG (new!)



Multipath TCP


extend TCP to allow **one** connection to transmit data along **multiple** paths between the **same** two end systems

pools capacity & reliability of multiple paths

traffic quickly moves away from congested paths

backwards compatible with TCP

Mon 17:40 MPTCP WG (new!)



App-layer traffic optimization

improve P2P performance

while at the same time aligning P2P traffic better with ISP constraints

provide P2P apps with network, topology & other info

enable P2P apps to do better-than-random initial peer selection

Wed 9:00 ALTO WG (new!)



Congestion exposure

explicitly reveal the expected congestion along an Internet path

new capability: allows even greater freedom over how capacity is shared

better info = better mechanisms

can be used for several purposes

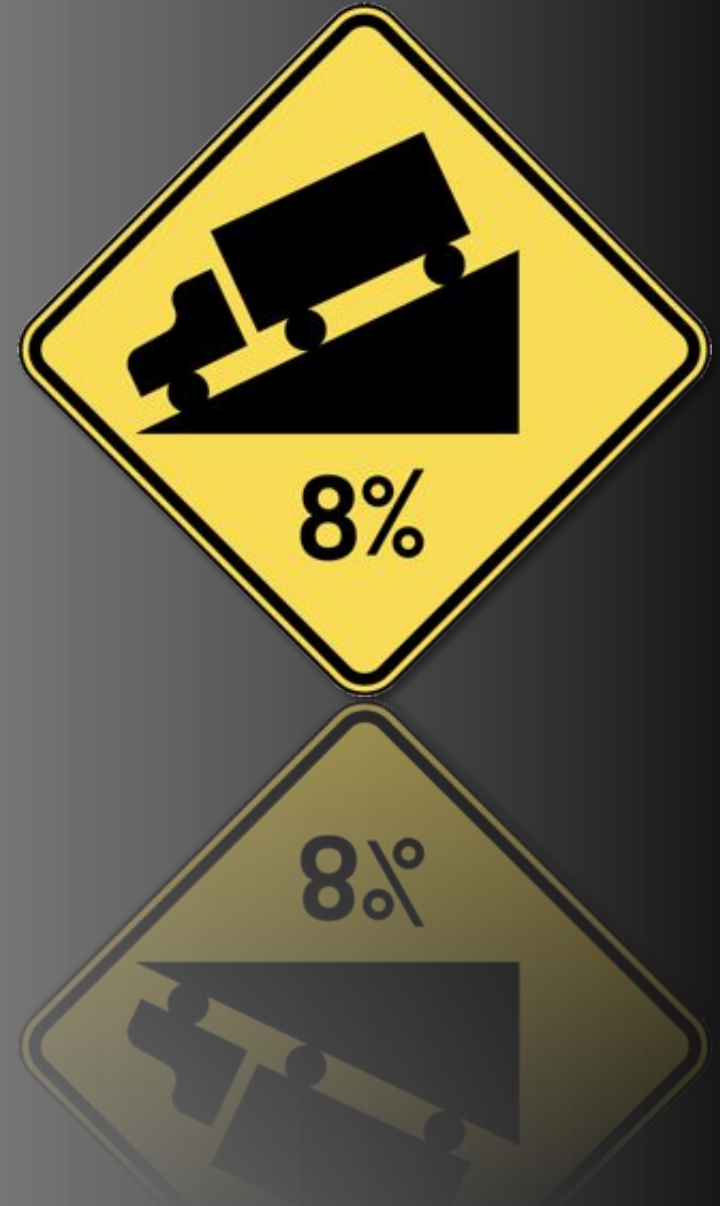
congestion policing

accountability

SLAs

traffic engineering

Tue 15:20 CONEX (BOF – not a WG)



Rec's for home gateways

goal: improve the network experience that a user of a home gateway gets when using the Internet

give an overview for implementers by collecting together requirements from different RFCs, e.g.,

- IPv4 & IPv6

- ECN & RED

- DNSSEC & DNS proxies

Mon 15:20 HOMEGATE (BOF)



**We have many tools to share capacity
fairly, effectively and efficiently.**

**And the IETF is designing new & better tools
where needed.**

**A lot could already be gained
by more consistently
and more appropriately using the tools we have.**