[li]nternet [applications] drivers of growth: 2005-2015

kc claffy
kc@caida.org
presented at
Telcordia Optical Networking Workshop
http://www.oida.org/
san francisco, ca
29 jun 2005

"How do you know when something is a 'killer app'? when every university tries to stop it and can't. that it takes a federal judge to threaten to put you in jail if you don't stop. that's how you know it's a killer app!" -- Eric Schmidt, 2001

outline of talk

- one slide about what was requested
  - application cross-section for next ten years
- why only one slide
- why it doesn't matter: economics
- implications for optical/fiber industry
- others who say this better than i do

measuring Internet traffic in the ‘core’

noone owns this problem.

- answer depends on network, no single answer
- providers don’t measure much themselves, due to enormous cost and low return (financial and other)
- safe generalities: http declining relative to p2p file sharing (often crossing eachother between 40% and 60%)
- email, IM, voice, filesharing here to stay, e.g., bittorrent > 50% on china telecom
- online gaming gets popular when bandwidth exists, e.g., south korea (bandwidth first, not demand)
- streaming negligible near as we can tell, no sign of changing
- why it is so hard to get a more accurate answer?: non-technical reasons block us first

apparently it's not as bad as i think

- don't panic (about that, anyway) because it doesn't matter: looks like this next decade is when we realize IP is like water.

- it's not food where you can charge more for making it look prettier. magic of IP is that endpoints can turn water into food. if you need to know application, you've already lost.

- you are not building laser-guided missles, you're building shovels. winners: stable, efficient, scalable, cheap, SIMPLE!

learning from [present] history

- where does innovation happen in the Internet?
- what does history consider failures? atm, rsvp, smart markets, active nets, multicasting, streaming real time multimedia, 3g
- what does history consider successes? email, web, browser, search, napster
- andrew odlyzko covers this much better than i

by some measures it does look bad

for U.S. inter-domain internet science, the crash happened in 1994 when the nsfnet retired...

- can’t figure out where an IP address is
- can't measure topology effectively in either direction, at any layer
- can't track propagation of a routing update across the Internet.
- can't get router to give you all available routes, just best routes
- can't get precise one-way delay from two places on the Internet
- can't get an hour of packets from the core
- can't get accurate flow counts from the core
- can't get anything from the core [we used to have anonymized traces]
- can't get topology of core
- can't get accurate bandwidth or capacity info
  - not even along a path, much less per link
- can’t trust whois registry data
- no general tool for `what's causing my problem now?’
- privacy/legal issues deter research (& it was hard in a enlightened monarchy)
- makes science challenging -- discouraging to academics

jarring observation from history of science

The modern field of elementary particle physics depended crucially on the establishment of a huge volume of data gathered mainly in the period 1945-65. Only then was it possible for the synthesis of the Standard Model to take place, 1967-74.

-- Peter Galison, Professor of History of Science and Physics, Harvard

(unfortunately, we're not doing research,
we're building critical infrastructure.
and it's riddled with structural problems.)

16 operational internet problems

- authentication
- security
- spam
- scalable configuration management
- robust scalability of routing system
- compromise of e2e principle
- dumb network
- measurement
- patch management
- “normal accidents”
- growth trends in traffic and user expectations
- time management and prioritization of tasks
- governance
- intellectual property and digital rights
- interdomain qos/emergency services
- inter-provider vendor/business coordination

(http://www.caida.org/outreach/presentations/2003/netproblems_lisa03)
(http://www.caida.org/outreach/presentations/2005/topproblemsnet)

why we're not making progress

- top unsolved problems in internet operations and engineering are rooted in **economics, ownership, and trust (EOT)**
- solutions require path "paved with incentives to motivate diverse agents to adopt it, implement it, use it, interface with it, or just tolerate it." without that path you've wasted your time.
- does not mean there are not useful technical problems to work on. but there will no technical solutions that don't solve the EOT problems.

warning: there's a problem we left out

- the economics one runs deep
- best available data suggests that moving IP packets around is not even a for-profit enterprise. not just bernie factor.
- like most large scale transport networks, it turns out. classic public good.
- noone tasked with thinking about the 25-year internet provisioning problem
- just as hard to get sound data on the economics

how unregulated players survive operating in an inherently non-profit industry

- hide the fact that you lose money by using non-IP revenue to subsidize developing IP habit. e.g., voice
- file bankruptcy every few years, includes billions spent on lobbying to keep incumbents in power rather than analyzing the macroproblem
- lie to the markets to get capital, confuse markets for a decade. or two. count on folks not reading history.
- long term: complex vertical integration (bad for security), infrastructure control (bad for freedom)
- don't let anyone look at the data that would facilitate analysis of provisioning models for this commodity. don't promote research & analysis.

what are we trying to do here?

The problem faced by the Internet industry is in ensuring that each provider of infrastructure is fairly compensated when its infrastructure is used. In essence, the problem is how to distribute the revenue gained from the retail sale of Internet access and services to the providers of carriage infrastructure. ....

While it is not completely clear that the deregulated open market nature of the Internet can sustain a diverse, efficient and effective service provider industry, it is also unclear what form of regulatory constraints or intervention are appropriate, if any.
.... There is a strong risk that regulatory involvement, if applied inappropriately, will trigger structural inefficiencies that ultimately will be reflected at the consumer level in higher prices and inferior services. Competition is not an end in itself, nor is regulatory impost. The challenge here is to foster the conditions that allow the Internet to be a productive and efficient platform for all. That, for me, appears to be at the heart of the challenge of the Information Society. Geoff Huston January 2005

(more succintly:)

“We never learned how to route money.” --David Clark, MIT

public vs private provision

- we don't really have a good grasp of the economics, social, or cultural impact, so aren't in a position to really say yet how the commodity of digital information transport should be best delivered to society
- govt has kept relative distance from Internet, for better and/or worse. changing.
- we guessed a lot of things wrong
- tremendous struggles for next few decades as we learn network economics the hard way (amidst multiple sources of measurement error)

implications for fiber industry

- reconceptualize internet as critical infrastructure
- that doesn't mean you won't make money (you will!), but on shovels, not missiles
- good news for you: foster new markets. (municipalities).
- health of your industry will depend on regulatory decisions more than applications

questions or corrections?
kc@caida.org