top problems of the Internet and how to help solve them

kc claffy
kc@caida.org

the significant problems we face cannot be solved by the same level of thinking that created them.
--Albert Einstein

16 problems of the net

- scalable configuration management (complexity, layering, legacy)
- security (aka 15 other things)
- host patching (aka ‘running bad software on purpose’)
- knowing what's on your network (measurement)
- spam
- authentication
- interdomain qos (aka ‘911’)
- compromise of e2e principle (aka ‘meeting market demand’)
- dumb network
- robust scalability of routing system
- “normal accidents” (charles perrow, we need him to study Internet)
- intellectual property and digital rights (aka ‘comatose industries’)
- governance (aka ‘regulation’)
- growth trends in traffic and user expectation
- inter-provider vendor/business coordination
- time management and prioritization of tasks

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

e.g.,Internet measurement: state-of-art

- can't measure topology effectively in either direction, at any layer
- can't track propagation of a BGP update across the Internet.
- can't get router to give you its whole RIB, just FIB (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 with real addresses in it
- can't get topology of core
- can't get accurate bandwidth or capacity info
  - not even along a path, much less per link
- SNMP just an albatross (enough to inspire telco envy)
- no 'why' tool: what's causing problem now?
- privacy/legal issues deter research
- how to build this missing theory? -- discouraging to academics

Result: measurements are a meager shadow of careening ecosystem.
[If you are not scared i am not explaining this right.]

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.

16 problems of the net

- scalable configuration management (complexity, layering, legacy)
- security (aka 15 other things)
- host patching (aka ‘running bad software on purpose’)
- knowing what's on your network (measurement)
- spam
- authentication
- interdomain qos (aka ‘911’)
- compromise of e2e principle (aka ‘meeting market demand’)
- dumb network
- robust scalability of routing system
- “normal accidents” (charles perrow, we need him to study Internet)
- intellectual property and digital rights (aka ‘comatose industries’)
- governance (aka ‘regulation’)
- growth trends in traffic and user expectation
- inter-provider vendor/business coordination
- time management and prioritization of tasks

why so persistently unsolvable?

- 1) rooted in non-technical issues: economics, ownership, and trust
- 2) not amenable to purely technical solutions -- require interdisciplinary investment
- 3) not amenable to < 4-year solutions (so even academia is out)
- 4) noone owns the problems

who can we get to help?

- ietf: lacks operational experience
- academia: lacks access to the infrastructure
  - no funding structure to incent attention to 20-year problems
- engineers (nanog): lacks (financial or legal) incentive structure to support cooperative work
- industry: lacks incentive
- content distribution industry: lacks basis in 21st century reality
- government: lacks clue regarding 21st century technology

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
- tremendous struggles for next few decades as we learn the economics the hard way (amidst multiple sources of measurement error)

pessimistic [near-term] reading

(aka “business as usual”)

- "the Internet is dying" -- Karl Auerbach provocative article
  - between spam, anti-spam blacklists, rogue packets, never-forgetting search engines, viruses,
  - old machines, bad regulatory bodies, and bad implementations
  - Internet will lose half its users in 6 months (yes that’s a bad thing)
  - in its place a much more controlled approved set of communications
  - lesson 1: run tcpdump if you don’t believe me -- most of it is garbage
- ”digital imprimateur" -- john walker
  - "how big brother and big media can put the Internet genie back in the bottle”
  - rich ‘optimistic pessimism”
- geoff huston’s nznog talk
  - trashing the Internet commons: implications for ISPs’
  - feb 2003, http://s2.r2.co.nz/20040129/ www.nznog.org/ghuston-trashing.pdf
  - not so much with the optimism
- bruce sterling keynote at NSF workshop feb 2002
  - http://www.cra.org/Activities/grand.challenges/sterling.html
  - ubicomp, ultrawideband, machines-building-machines are his messiah

optimistic [long-term] reading

(aka “living up to the net’s potential”)

- eben moglen, columbia
  - freedom of thought
- lawrence lessig, stanford
  - code is law, future of ideas
- yochai benkler, yale
  - linux and the nature of the firm

constituency responses to situation

- federally funded research community: battles clock
- nsf: battles incrementalism
- operational community: battles worms & growth w/ minimal cooperation
- telecom: battles antique regulation, unprofitability/mergers, we-the-people
- copyright-owners: battles its own customers with state backing
- fcc: battles irrelevancy (and wardrobe malfunctions, howard stern)
- military: battlefield ISP (profoundly important)
- people who can afford it: build their own (google, aol, DOE, cenic)
- rest of world: battles digital divide, us hegemony
- UN/ITU: grasps for power

implications for empirical Internet research

- need to start asking questions we ask of critical infrastructure
- need vehicles to inform policy
- need vehicles to protect & analyze data

implications for architectural Internet research

- goal: design `in the light' (first time ever..)
- need interdisciplinary, multi-agency investment
- need longer-term thinking than any current vehicle supports, e.g., need to anticipate not just technical but social and political trends 25 years out
- non-hierarchical frameworks merit attention

implications for intellectual property

- assumption: in 25 years, everyone has unmediated connectivity to everyone else. (political certainty: all alternatives are worse.)
- stronger copyright protection for cyberspace: “desirable, inevitable, and irrelevant.” (andrew odlyzko said 9 years ago). code is law.
- cost of distribution -> 0. companies who charge for free services will go away

implications for regulatory research

- goal: bring regulation and economic models in congruity with technology and empirical data
- investigate alternative models of provisioning
- CENIC-like organizations can play a role here
- so can estonia
  - http://www.privacyinternational.org/survey/phr2003/countries/estonia.htm "The 1992 Estonia Constitution recognizes the right of privacy, secrecy of communications, and data protection."

implications for CENIC

- hopefully you'll be inspired to help
- push your fiscal and other lessons outside the CENIC community into your communities, to the state, to the world
- enlightened connectivity models will serve 'fittest' societies in the 21st centuries
- north star: most economic way to promote freedom of access to all
- promote open analyses of cost models: how much is unfettered p2p digital access worth?

should california have a cenic?

- a la stem cell initiative
- teach people that provisioning models can change as technologies and goals change
- proposition 215 for spectrum in california?
  - using our best understanding of the world, narrow the congruity gap between policy and science
  - demonstrate open spectrum has higher social value than cost
  - should we ignore federal stance and just start using our air? is leading by example the best way toward informed, enlightened policy?

the battle, sir, is not to the strong alone;
it is to the vigilant, the active, the brave.
-- patrick henry

kc@caida.org
www.caida.org