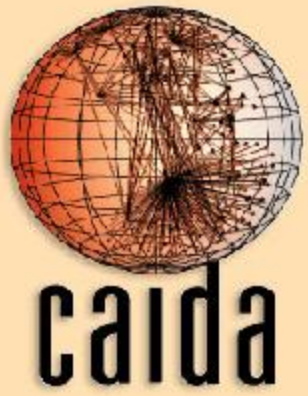

priorities and challenges in Internet measurement, simulation, and analysis

problems that remain persistently insolvable
should always be suspected as
questions asked in the wrong way.
-- alan watts

10 june 2003
lsn meeting
kc

objective of this talk (per LSN request)

identify R&D gaps

- large-scale deployment issues for federal agencies that fund network research
  - DOE, NSF, DARPA, NASA, NSA, NIST

themes

- develop measurement as a respected field of science and engineering
- address relationship/rift between research and operations
- acknowledge problems persistently unsolvable in current paradigm
- in the face of brilliant expectations, demand, and opportunities, the need for network self-awareness is crossing a threshold for both the high and low end
- recession combined w/heightened consciousness of homeland security creates an inflection point in the opportunity for U.S. policymakers to make a positive difference

outline of this talk

- why Internet provision market forces are torqued
  - how to make the problem worse
  - effect on research community
- why policy makers have a vital role to play now more than ever
- example of perfectly normal accident
  - why it matters
- long-term challenges in Internet measurement
  - 12-step program as outlined by NSF ANRI PI meeting measurement breakout group
- near-term community gaps in measurement
  - concrete 1-2 year tasks
- action items for research community (slide 24)
- action items for funding agencies (slide 25)

not our father's Internet

radical changes in last 20 years

- capacities
  - increased by 5+ orders of magnitude
- penetration
  - universities -> worldwide commercial and residential
  - 170M+ hosts
- usage
  - academic -> commerce, government, entertainment, porn, spam, broken traffic
  - not necessarily in that order
- trust models
  - cream cheese -> swiss cheese
- service models
  - best effort -> VoIP, e2e model, VPN
  - except w.o definitions
- e2e architecture
  - guiding architectural principle -> historical artifact
  - still the most disputed of Internet holy ground

accidental funding model of Internet

more radical changes

- originally: funding was for technology not infrastructure
- NSF took up infrastructural gauntlet for R&E community
- put it down when it was already too hot to touch

- built and deployed architecture with inability to account for costs
  - with any degree of granlarity
- inconvenient when we finally needed someone to pay for it

result: market forces are badly torqued

result: market forces torqued

R&E casualties of Internet industrial (!r)evolution

- spam
- multicast
- qos
- DOS attacks
- routing announcements
- address space
- UDP, unfriendly TCP
- death of e2e principle
- software bugs
  - _normal accidents_ (charles perrow)
- measurement (meta-issue, integral to all of the above)

worse than `tragedy of the commons': active disincentives to fix

- competitive pressure, pricing models, privacy, constitutionality, technology, lack of standards, dearth of capital

ways to make the problem worse

- make network management research seem really boring
- lend measurement no respect as a field in and of itself
  - don't fund it
- don't require any real data from providers
- let market forces take care of it

if any free market has proven incapable of taking care of itself, the Internet is it.

if any system has proven itself more expensive to not measure than to measure, the Internet is it

- see IAB recommendations regarding Internet research & evolution
- see global recession

what this problem is not

- not a public park
- not libraries
- not the phone system
- not the electricity grid
- not the railroad
- not the highway system
- not tragedy of the commons
- not cathedral, not bazaar

these analogies all break down. this is something new.

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

inherent operating constraints of constituents

- vendors
  - provide minimum measurement functionality (snmp, netflow)
  - no incentive to do more since it is expensive and does not sell boxes
- providers
  - dis-incented to share data with others due to competitive and privacy concerns
  - no apparent payoff (beyond what they do already)
- software engineering (operating systems, applications, routers)
  - lack Internet systems perspective
  - renders Internet vulnerable or non-optimal
- measurers/researchers
  - oblivious to real world of costs, e.g. opex/capex ratios
  - largley unable to target measurements at immediate provider problems
  - struggle with interpretation of non-standard data sets
- users
  - as if we could forget

if policy makers ever had a role to play anywhere..

(manipulator of market forces to protect/enhance consumer welfare)

unsettling admissions about dealing w data

[courtesy vern paxson & david moore:]

- www.icir.org/vern/talks/vp-nrdm01.ps.gz
- www.caida.org/outreach/presentations/2002/ipam0203/

- bizarre behavior, misconfigurations, non-RFC, attacks, `impossible' behavior
- measurement tools lie
  - (packet filters drop, reorder, replicate, miss due to routing)
  - clocks can be arbitrarily off/moving, timestamps don't know accuracy, applied differently
  - app-level measurement tools miss hidden network stuff (middleware, socket buffer parameters)
- asymmetric paths
- measurements made two different ways always disagree (anisotropic)
- even a single measurement may disagree (not atomic): routing tables, traceroute
- events ripple through network along trajectory that is unlikely fully instrumented
- measurements carry no indication of quality
- measurements lack meta-info (e.g., hostnames)
- representative data points - there is no typical on the Internet
- analysis results not reproducible
- we lack a culture of calibration
- large-scale measurements required for representative/longitudinal analysis overwhelm our current methods
  - archived data often ad hoc, corrupt, truncated, poorly documented. unnavigable
  - lack of historical data renders it difficult to assess trends
  - alas, people do it anyway, see kc's myths talk (or any trade rag)
- Internet measurement, although too hard, is too easy
  - not enough data and too much data
- we don't yet know how to measure real traffic in the core
  - speed, sampling, anonymization
- can't keep up with media in core (oc12 monitor arrives right after upgrade to oc48)

in spite of it all, amazing stuff has emerged

recession has helped raise awareness of the value of data

(see nanog 2002-2003 meetings, nanog.org videoarchives)

DNS:

- damage in the DNS system

BGP:

- route dampening considered harmful (again)
- behavior under stress (lixia et al)
- non-deterministic routing can be demonstrated w/in full mesh topology
  - enhancements to prevent persistent route oscillations
  - partial order (due to MED) in route selection
- MEDs prevent reasoning about routing policies
  - though existing references encourage use of MEDs to influence inbound policies
- how much are we facing limitations of distance vector protocol and inherent limitations of same

lots more examples

upshot: *environment more receptive to collaboration than ever*

normal accidents: rfc1918 traffic sent to roots

updates for private addresses leak outside local domains

- spectroscopy analysis of RFC1918 updates coming from DHCP/nameservers
- tens of millions a day coming to blackhole servers
  - **51.4M updates in 86.5 hrs = 10,000/min = 165/s. up to 1200 updates/s (nov02), up to 29 pkts/update**

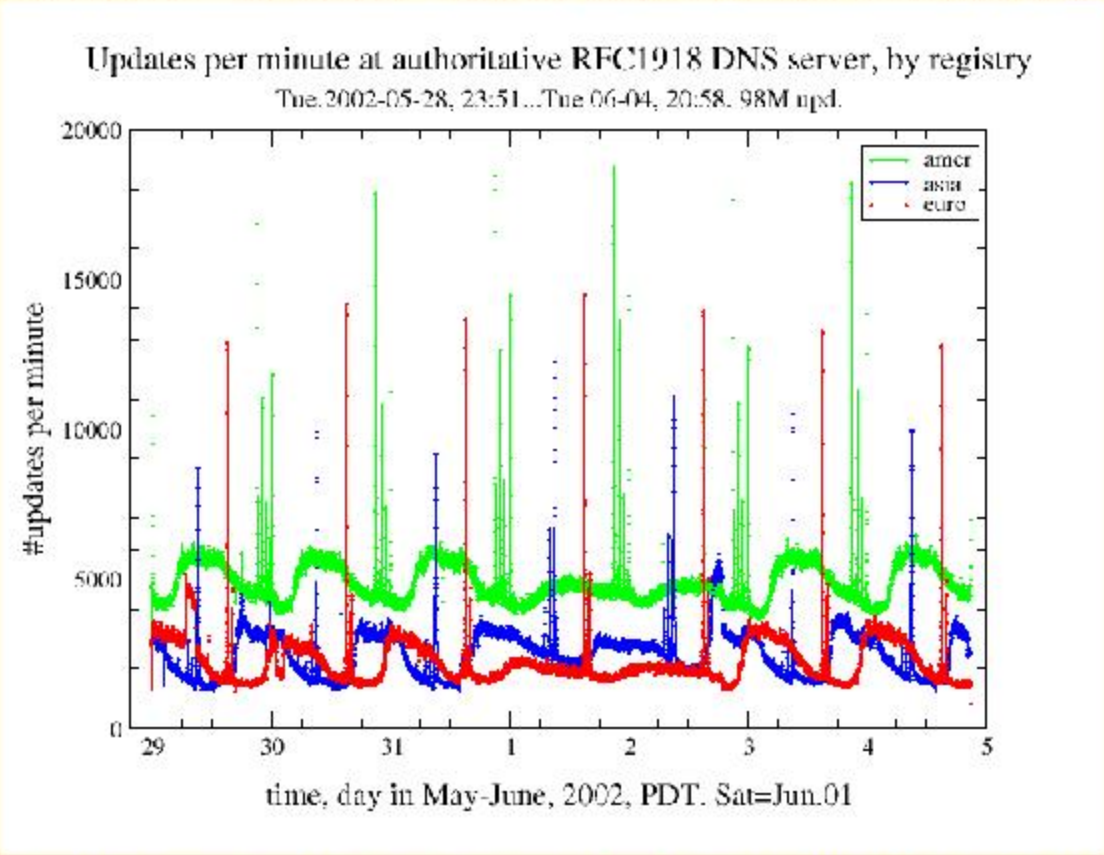

- **weekday, weekend patterns; weird spikes at midnight local time (4 in US, 3 in Asia, 2 in Europe)**
- **can see that Asians work on the weekend, Europeans not so much**
- **can see that Europeans and Asians get to work on time**

... global RFC1918 damage in DNS system (2)

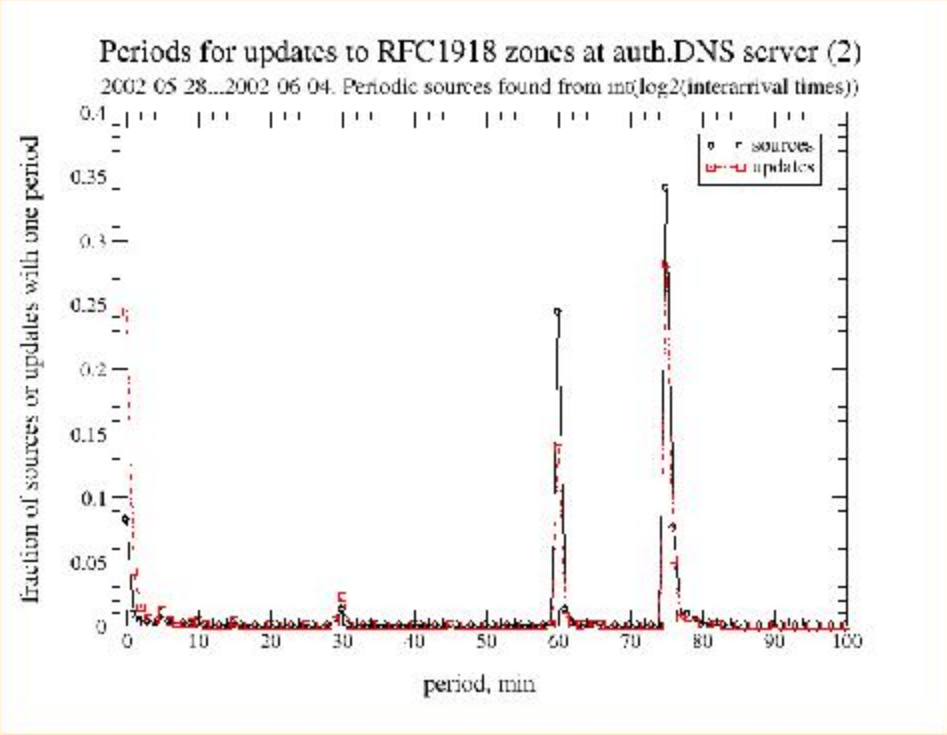

- rare to get macroscopic Internet data so radically broken
- who is trying to update the roots anyway?
  - **dsl, cablemodem, small population providers, developing countries**
- verified that vast majority derive from two OSes: Windows 2000 and Windows XP
  - **majority of updates from sources that send them constantly**
  - **bulk of workload from contributions of medium size, not mice/elephants**
  - **most source IP addresses are of home and small business users (owned by individuals, not organizations)**
  - **connected to the Internet via cable, DSL or phone-based ISPs**
  - **majority using software with default vendor settings**
  - **academic, corporate, backbone networks contribute little rfc1918 update traffic**

...global threat arising from single vendor

- combination of Microsoft software features & misconfigurations was essentially causing a slowly paced massive distributed denial of service (DDOS) attack on the root name server system
- current state of fielded desktop software poses substantial & increasing burden on (if not threat to) the robustness of the global Internet
- software and setups affecting global systemic Internet stability must be designed more carefully wrt potential effects of:
  - software engineering decisions
  - misimplementations
  - misconfigurations
- measurement can make a huge difference

optimism (no, really)

still the case: security, performance, configuration, and fault management lack both effective solutions, as well as an apparent lack of people able to state a concise problem to be solved.

[ok so that wasn't the optimistic part...]

researchers do have a valuable role to play including because we are a(n albeit relatively) trusted neutral party.

but we have to explore outside of our labs.

challenges in Internet measurement (1)

(12-step program, from NSF ANIR PI meeting breakout session, jan 2003)

motivating vision: self-aware network

- cultivate culture of sound measurement as science & discipline
  - measurements need pedigrees describing them, how to navigate
  - audit trails, portable analysis scripting language to support reproducibility
  - well-managed meta-data
  - understand sampling implications and technology better
  - anonymization tools & reduction agents
- more strategic measurement, guided by rather than constraining research questions
  - what data is missing and how do we strategically optimize the return on investment in data collection and instrumentation
  - recognizing that we don't always know what questions will be asked next year
  - improved standardized interface to data archives

challenges in Internet simulation & modeling(2)

- mathematical frameworks to find structure/patterns in traffic
  - a la scott's encouragement to `formalize some of what we (and providers) know'
  - macroscopic as well as microscopic
  - theory of joint spatial/temporal locality
  - spectroscopy, tomography
  - ietf/ippm has been trying for a few years, but without dedicated funding
- source modeling (for realistic inputs into simulations, models)
  - extract a set of source models from an aggregate trace
    - **feature extraction problem**
    - **10,000 gnutella port numbers are not 10,000 flows**
  - ultimate goal: augment libraries of source level models w generation of own
  - calibrate models by evaluating their power for prediction
- empirically validated simulation of a significant aspect of the Internet
  - already much work in large-scale simulations, but no recognized empirically validated simulation of any signficant piece of the Internet
    - **note: large scale means in size as well as # of protocols**
  - requires cooperation from providers and vendors to get default and configured parameters of OSes and algorithms. govt could shepherd/foster this cooperation

challenges in Internet analysis (3)

don't forget the real world

- analyses must incorporate expense ratios (opex/capex) into tradeoffs where possible
  - need to get/keep relevant to providers
- systematic studies of outages
  - assessment of various causes of damage
- tools and techniques to tie user-perceived performance w/ network measurement, with statistically significant results
  - function of control plane, routing plane, server on other end.
    - **including bugs in routers, servers, client stacks**
  - correlating user experience with events that happen in the network
  - incorporate into network performance models
  - validation of tools still doesn't exist
  - use at truly large scale still doesn't exist
  - dave clark's `why?' tool (non network geek compatible) still doesn't exist

other challenges in Internet measurement (4)

- discovering pervasive hidden bugs
  - any modeling or analylsis must also handle the impact of this huge component of traffic
- how does measurement affect/support security goals
  - infer bgp, firewall, and virus spread behavior
  - how do you get networks to share security-related information
  - protection of measurement infrastructure from security compromises
- measurement specific to optical, wireless and sensor networks
  - especially assessment of new application domains (for wireless/sensor)
- encouragement of strategic measurement in new networks
  - based on what we learned from what we did wrong in old networks

near-term community gaps in measurement (1)

(high impact tasks for next 1-2 years)

problem diagnosis & response

- detection, location, isolation & reporting
- macroscopically, DHS (homeland security)'s GEWIS problem
  - global early warning information system
  - DNS
  - BGP
  - not so near term, needs support from <next 4 slides>

performance

- low impact bandwidth estimation
  - ideally single source
  - scaling to Gbits/s & new NICs (interrupt coalescing)
  - calibration against real paths w cross-traffic
- e2e trouble-shooting toolkit
  - 'why?' tool
- SLA validation
- large scale RTT distributions, publically available

near-term community gaps in measurement (2)

security

- traceback, forensics
- network telescope (backscatter analysis)
- automated worm response
- intrusion detection

topology/routing

- `pop/router level map of the Internet'
- AS connectivity ranking
- refined methodology for topology and routing measurement
- strategically designed interdomain routing data collection
  - constant, sufficiently large, diverse, and representative coverage
  - stability across time
  - IGP updates (requires cooperation/release from ISPs)
  - configured parameters from ISPs, e.g MRAI, metrics
- tracking topology at other layers
- IPv6
- aggregation, abstraction, and visualization techniques

near-term community gaps in measurement (3)

traffic characterization

- workload and flow analysis and modeling
  - longer traces (at least 24 hours, diverse sites, times)
- longitudinal trend analysis (baseline)
- router/switch capabilites
- options at range of cost/capability

wireless, sensor

- headers, location, signal strength
- correlation among data from massive nodes

community gaps in measurement (4)

meta-problems

- firewalls, filtering, blocking (ports/apps)
- VPNs, layer2
- archiving
- privacy/sharing
- interpolation/extrapolation
- validation
- correlation
- feedback to system at various granularities (dynamics and trends)
- new protocol development: routing, transport, IP
- feedback to future operational measurement methodologies

action items for research community

- culture of & passion for sound measurement, as science & discipline
  - measurements need pedigrees describing them, how to navigate
  - audit trails, portable analysis scripting language to support reproducibility
  - well-managed meta-data
  - understand sampling implications and technology better
  - anonymization tools & reduction agents
- simulation
  - need way to calibrate against real data (broken record)
  - safety tip: still has mostly no respect from providers. diplomacy serves.
- analysis
  - find ways to assess opex versus capex of any new idea
    - **or at least don't render it impossible to do so later**
  - as scott encourages 'let's try to formalize some of what we (and providers) know'
  - as dave clark encourages 'it's about the $$$, stupid'
  - safety tip: providers do often quickly lose their patience w researchers
- continued/increased interaction with providers and vendors
  - nanog, ietf, caida
    - **ask for help from those who have succeeded**
  - switch, router, measurement hardware vendors show more interest than ever
    - **people don't trust `analyst estimates' anymore... good news in many ways**

take advantage of industry regrouping efforts & inclination to listen

action items for funding/publishing agents

- set measurement standards for procurement of federal agency networking services (qos, security)
  - inter-agency measurement/qos agenda (start w cross-community workshop, might result in dedicated FTEs at each network..)
  - new LSN working group, perhaps eventually program, focused on measurement
- willingness to invest $$$$ in measurement and databases
  - software, storage, communications for collaboration, data distribution
  - frustrating [perception] that it takes funding away from `real' research
  - encourage tech transfer from other disciplines who have done it
- willingness to bounce manuscripts that don't include raw data + scripts
- willingness to publish `mundane' work on measurement management
- emphasis on reproducible results
- encourage research projects that involve provider/vendor cooperation
  - and discourages ones that don't
- encourage researchers to participate in industry meetings and mailing lists (nanog, ietf)

good news: there has never been a better time for this

- overheard 'the recession is a stay of execution for the bgp routing system'

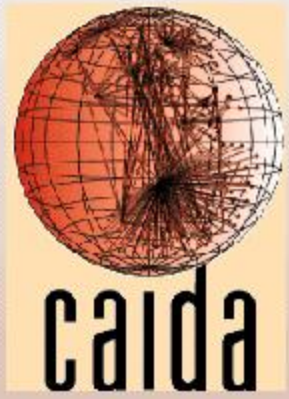

// disorder increases with time
because we measure time in the
direction in which disorder increases //
-- stephen hawking

www.caida.org