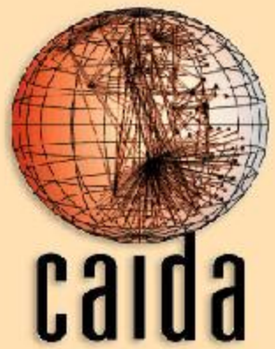

top problems of the Internet and what sysadmins and researchers can do to help

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

october 2003
kc

outline

- what talk(s) i'm not going to give
  - lofty things like telepresence, virtual reality, ubicomp, emergency response system
  - specific measurement & modeling priorities (see caida web site for that talk)
- what NSF (and other funding agencies) consider top problems
- what problems i will cover
  - those more relevant to sysadmin audience
  - either because you can help or because you need to stay cognizant
  - 1-2 year not 5-10 years out (am a big fan of far out but won't emphasize it today)
- underlying themes: `interactive complexity' and `space in between'
  - charles perrow's `normal accidents'
- list of top problems
  - i thought i should at least mention them
- "sky is falling" vignettes
  - reading recommendations
- why system and network administrators are key to solutions
  - sysadmins: if you manage systems connected to the Internet, you're a network administrator
  - add it to your business card so you remember

acknowledgments/caveats

this talk in the `if you steal from one author, it's plagiarism; if you steal from many, it's research' category

dozens of conversations with variety of experts

- could kibitz for another 6 months but answers would likely change
  - actually that's ridiculously optimistic
- wrote down all contributing names so have them in the raw data
- but those who spent hours discussing it
  - mike lloyd and sean finn (routescience)
  - johna johnson (nemertes)
  - some iab, iesg, ietf folk
  - funding agents
  - network engineers, security folk
  - Internet researchers
- didn't ask (to my knowledge)
  - spammers, hackers, layer1 folk, lawyers, teenagers, K12 teachers, librarians, FBI/CIA/NSA, RIAA, MPAA, john ashcroft
  - so i don't claim to represent them

occasional theme of talk

chick flick before sunrise's *"the space in between"*

system and network administrators are
the ultimate 'space in between' (two user nodes)

- exalted sense of that phrase
- often the only intelligent glue holding two users together
  - we oversignificate protons and neutrons too, maybe because we understand them better than we understand the space in between
- critically important but underrepresented in policy, protocol, architectural areas

confession: if this talk comes across as a call to arms,
i can live with that.

a view from National Science Foundation

NSF workshop on fundamental research in networking

- april 2003 workshop
  - all position papers online (go nsf!)
- federal community takes prioritization of goals pretty seriously (yay!)
  - goals may not always be prescient, or even sensible
  - but it's not because of indifference
  - some of these people are damn good
- will read some (not all) recommendations from report
  - listen for your name!
- then discuss innovations needed to support these grand challenges
  - cart is now officially before horse
  - essential pieces are being confidently assumed from both directions

NSF recommendations in april 2003 report

reflections on past, present, & future of networking research suggest need for emphasis on

- radical innovation and paradigm shifts
- multi-disciplinary aspects of fundamental research in networking
- outside-the-box thinking, beyond the success of the Internet
  - avoid "network innovator's dilemma" -- too involved in improving the existing Internet technology to observe novel and disruptive technologies
- increased reproducibility of networking research
  - ease burden of reproducible experiments on complex systems

strong sense in here of `being done with the Internet'
as if the most important research problems there have been solved

NSF recommendations in april 2003 report

and yet... grand challenge recommendations

- focus research agenda on application-induced challenges
  - user-focused
    - **robustness**
    - **transparence, ease of configuration**
    - **exploiting storage and processing capabilities**
    - **power consumption**
  - network-centric
    - **heterogeneity and scale**
    - **manageability**
    - **evolvability**
- meta challenge: in face of shortcomings of current Internet architecture, develop new network theories, architectures, and methodologies that will facilitate the development and deployment of the next generation of services and applications

undeniable (tho not verbatim) admission that shortcomings of current architecture all relate to the inability to manage (administer) it

the plot thickens...where are folks we expect to manage it?

grand challenges: highlights of two i care about

(still from nsf workshop report april 2003)

- 1) Internet information theory
  - holy grail: Internet Erlang
  - for wireless too
- 2) overlay networks
  - all getting along, even economically
- 3) network economic theory (network markets)
  - whole slide on this one (next)
- 4) resilient networking
  - whole slide on this one (next+1)
- 5) sensorized universe
  - security, health, education, battlefield, traffic, law
  - holy grail: `ubiquitous safety net' (emergency response system)
  - less holy grail: bigbrothernet
- 6) virtual networks
  - VPNs, mobility, automatic team discovery/creation, cognitive networking

challenge (3) network economic theory

lack thereof oft attributed cause for failure of qos, multicast, CDNs, pick an Internet startup space.

- difficult to make progress with such an uninstrumented network
- micropayments have failed a few times but will likely get another chance
- architectural bottleneck often considered global network economics rather than technology

- current economic model fails to capture potential utility of network
- so Internet can't support many potentially cool applications
- need to get some economists in here

challenge (4) resilient networking

faults are norm rather than the exception

- component failures, human operational errors, software viruses, malicious attack

NSF-articulated needs to support this challenge

- system development tools that reduce frequency & severity of bugs
- programming languages, environments, audit tools, runtime checking tools that reduce frequency & severity of config errors
- understandable, deployable, and usable security
- new approaches to the composition of modular elements
- new approaches to federation
- pervasive audit trails
- self-adaptive systems (automatic detection of and response to attacks, route changes, errors)

that looks dangerously close to something we need now

(sysadmins: if you don't see your name above you're not paying attention)

meta-challenge: conquering system complexity

related to resiliency

- going to go on a **normal accidents** digression now
- yale sociologist charles perrow book, 1999 (latest edition)
- focus on large scale systems more tightly coupled than ever imagined
- *interactive complexity*, coupling, and catastrophic potential
- Internet not mentioned once in this book
  - but you'd never believe this guy hadn't configured a BGP session

why i hope dr perrow writes a book on the Internet next:

(from same NSF workshop, http://www.cra.org/reports/gc.systems.pdf)

"rule of thumb in production software about 70% of the code deals with error conditions and exceptions, and only the remaining 30% provides the functionality expected by users. In complex systems with billions of components, the portion of code devoted to functionality will decrease further; most code will deal with adaptation to dynamic change, not just error handling, but overall self-sustainment."

normal accidents (perrow, 1999)

theme of book

- way the parts fit together, interact, that is important
- the dangerous accidents lie in the system, not in the components
  - the 'space in between'
- air transport system works well
  - diverse interests and technological changes support one another
- all human constructions resistant to change
  - driven by private privileges and profit
- fairly optimistic ending
  - (well depending on your perspective)
  - Y2K as analogy: huge global socio cultural momentum does make a difference
  - apparently we don't have an 'imminent common globally destabilizing threat' right now
    - **noone's more surprised than i**

normal accidents: early excerpt

``for some systems that have this kind of complexity, such as
universities or research and development labs, the accident
will not spread and be serious because there is a lot of slack
available, and time to spare, and other ways to get things done.

*[kc: above was written in the mid 80s and doesn't describe many universities today.
nothing on the Internet can safely assume `time to spare'.]*

but suppose the system is also 'tightly coupled', that is,
processes happen very fast and can't be turned off,
the failed parts cannot be isolated from other parts, or
there is no other way to keep the production going safely.
then recovery from the initial disturbance is not possible;
it will spread quickly and irretrievably for at least some time.
indeed, operator action or the safety systems may make it worse,
since for a time it is not known what the problem really is."

-- p.5 normal accidents

if that doesn't remind you of debugging BGP configs it's because you haven't.

- most BGP admins i know are more "used to" than "understanding" configuration
- too many people are configuring BGP for us to accept any less transparency
  than with driving a car
- ok maybe a truck. but definitely not a space shuttle. we abjectly lack a houston.

normal accidents (perrow 1999)

complexity, coupling, and catastrophic potential

- inspired by accidents like TMI, chemical plant mishaps, aircraft collisions
  - also a postscript on Y2K written in 1999
- compelling delineation of four classes of victims affected by an accident
  - (1) operators (2) users (3) system outsiders (4) future (fetus, future user)
- passive versus active risks
  - passive -- safety beyond user control
    - **airline, concerts, shopping malls**
    - **generally someone else making a profit.**
  - using Internet is passive risk for most people
  - but can be passive as in conscious but not intentional risk ('price of convenience')
  - but on the Internet many folks not even conscious of the risks they take
    - **as if sysadmins need to be told this**

normal accidents: how to minimize

daunting to track safety of pervasive complex system

- safety records don't normalize for number or expertise of participants
- as `user-friendly' technology reduces risk, more users join activity
- but the clueful ones are already on, so end result is that
  even as component safety increases, the accident level may not change!

the safer you make it, the lower clue threshold needed to operate, the safer you have to make it

(perhaps i don't need to dwell here since if anyone knows that
trying to compete with stupidity is a losing battle,
it's system administrators)

- related dialectic between Internet robustness and complexity
  - see http://www.1-4-5.net/~dmm/complexity_and_the_internet/

normal accidents: garbage can theory

'garbage-can theory' (bounded rationality)

- describes decision-making in highly ambiguous settings
- "organized anarchies"
- irregular confluence of people, problems, solutions
- uncertain circumstances, and choice opportunities
- decision makers move from one opportunity to the other
- relying on chance alignment of components and organizational demands

- published in field of organizational behavior
  - Cohen March and Olsen (1972)
- reckon network research has similar but we have studied naming a lot
  - pending funding for 'Internet garbage can theory'
  - in meantime we also have dilbert as mouthpiece for basic tenets
  - actually we need more Internet garbage can science
  - since we have not formalized a lot of our garbage yet
  - see dave plonka's lisa talk, caida rfc1918 paper, redundant anycast traffic, etc
  - grep for garbage in bruce sterlings's nsf keynote talk
    - **http://www.cra.org/Activities/grand.challenges/sterling.html**
    - **exceptionally worth reading anyway**

normal accidents: closing quote

catastrophes send us warning signals. this book
has attempted to decode these signals: abandon this,
it is beyond your capabilities; redesign this, regardless
of short-run costs; regulate this, regardless of the
imperfections of regulation. but like the operators of TMI
[three-mile island] who could not conceive of the worst --
and thus could not see the disasters facing them -- we
have misread these signals too often, reinterpreting them
to fit our preconceptions. better training alone will not
solve the problem, or promise that it won't happen again.
worse yet, we may accept the preconception that military
superiority and private profits are worth the risks.
this book's decoding asserts that the problems are not
with individual motives, individual errors, or even
political ideologies. the signals come from systems,
technological, and economic. they are systems that elites
have constructed, and thus can be changed or abandoned.

--normal accidents, charles perrow, 1999

the part of the talk you came for

top problems of the Internet

compiled from dozens of the most brilliant people i know

problem of the Internet

scalable configuration management

- higher layer connectivity requirements hard are to express, manage, maintain, verify still working, simulate, model
- today's routing configuration languages are based on low-level mechanism, rather than operator intent
- networks are configured at the element (or router) level, rather than as a single cohesive unit with well-defined policies and constraints
  - key network operations goals require tweaking configs in pursuit of desired indirect effect on the network
    - **traffic engineering, security**
- usual mode of coping: monitor for things that break
  - not things that -might- break if you make a change
  - use Internet as a simulator
    - **"`current best practices?' is that a band?"**
- lots of things to configure, even along one path:
  - router, switch, load balancer, (NAT) host, OS, web server, application, database.
- configuration management is everywhere
- word for the decade?: abstraction

scalable configuration management (2)

partially responsible for current situation

- trusting vendor defaults too much
- putting up with vendor kitsch
- no trusted routing registry
- egregious lack of instrumentation
- moderate lack of clue
- business constraints retrohacked into a system not designed for it
  - garbage can theory!
- inherent interactive complexity of global distributed system
- interdomain (BGP) routing system configuration gets its own slides later

what can be done

- researchers: "higher level policy languages"
  - abstraction abstraction abstraction
- sysadmins: help define scope of your configuration needs
  - am not claiming we can definitely get there with incremental steps
  - only that we don't have an alternative at the moment

problem of the Internet

security

- also known as authentication, availability, containment, DOS tracking, identification, privacy, robustness, resliency, recovery, & threat analysis
- could include spam depending on party (we'll award spam separately)
- solutions to vastly [un]defined problems are inherently elusive
- something we have learned for certain: cryptographic algorithms and standards for authentication, security, and privacy are far ahead of our ability to deploy, administer, and use security systems
- need new specification techniques for security policies
  - meaningful to system administrators and end-users
  - so security is deployed in a way that meets user expectations
- increased automation, rigorous analysis, baseline profiling data
- self-configurable and self-healing systems
- ISP cooperation (DOS traceback)

sysadmins: if you think these problems will be solved by another community
i encourage you to investigate further
because whoever is solving them needs your help anyway

problem of the Internet

end host patching

- better patch clue
  - patches can make problem worse, break other things
  - if a patch does that, please tell your vendor...
    - **example: code red -- people couldn't patch IIS without breaking realsecure so many didn't patch**
- 'default deny' is your friend -- at host level!
- help develop or at least be aware of product liability laws
- won't push genetic diversity argument as alternative `safety'
  - sounds too security through obscurity to me
  - unclear how much manageability would be sacrificed to get it
    - **already too much whack-a-mole in this field**
    - **fidelity.com (who handles about a billion dollars a day on the Internet) already can't handle my mozilla**
    - **if we espouse genetic diversity, we better espouse a hell of a lot of systemic investment in software testing**
  - besides hey i'd run a monopoly OS too were it the best OS
    - **monoculture paper suggests it might not be possible: http://www.ccianet.org/papers/cyberinsecurity.pdf**
  - many unixes use RPC and same BSD stack anyway
  - most importantly, it may be a good idea but it is no substitute for patch clue
  - illegitimate botnetting is big financially backed industry now
    - **serious income motivation to find holes**
    - **see rob thomas' aerobic nanog talk online (oct 2003 meeting)**
  - a few more OSes on the Internet would not diminish the catastrophic potential
  - the kiddie scripts would just be longer

problem of the Internet

knowing what's on your network

- how many site administrators run one of:
  - flowscan, flowtools, netflow, autofocus
  - see tool taxonomy http://www.caida.org/tools/taxonomy/workload/
  - latest addition: UCSD CSE's http://www.caida.org/tools/measurement/autofocus/
- follow relevant R&D measurement activities and peer-reviewed tools
  - IETF WGs, e.g, IPFIX NANOG, sigcom, IMC, PAM
  - work with researchers on tools and visualization techniques

measurement is such a dead giveaway, it advances ball on

- capacity engineering
- security
- privacy
  - indirectly... teach users the realities of measurement. then teach them ssh and pgp
- provider integrity checks
  - the more grassroots measurements, less likelihood of another irrational bubble
- obligatory caveat: know the law
  - need measurement tools that help manage & secure your network without breaking the law
  - we need your help getting better laws

know what's on your network (2)

not to imply that measuring the Internet in general works

- can't measure topology effectively in either direction. at any layer.
- can't track propagation of a bgp update across the Internet
  - so how to build this theory we're so lame for not having -- discouraging to academics
- 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 disincent research
- result --> meager shadow of careening ecosystem

if you're not scared i'm not explaining this right

problem of the Internet

spam

- consider this more of a user issue than an Internet issue
- we are relying on some ad-hoc defacto messaging systems (SMTP, IM) that were never designed for corporate high-integrity use
  - you aren't going to stop spam with this version of SMTP
  - need new messaging infrastructure with built-in authentication
  - some work going on in IETF; please participate if you care at that level
- in meantime current network-level cures are worse than the disease
  - an ISP or sysadmin blocking traffic for content is a dangerously slippery slope
  - not clear we want to live in that kind of world
  - yet another arms race (along w p2p, firewalls, verisign wildcards)
  - you're in position to mention to those who might have other interests driving their behavior

what can be done to help

- give your users plenty of client side options for filtering
- give operationally flavored input to IETF activities in this area
- adjust expectations (not to be confused with admitting defeat)
  - advertising has been no enemy to our free (not to mention cheap) press

problem of the Internet

authentication

- mentioned earlier under security and a second ago with spam
- also often called `the identity problem'
- not to be confused with 'anonymity' which isn't on our problem list
  - like spam, more of a user issue, should solve outside the architecture
- lack scalable, non-hierarchical trust models

<rant>

btw putting a "solution" label on pgp is ludicrous

- pgp(/gpg) remains an egregious tech transfer failure
  - perfect example of algorithms/standards far outpacing ability to deploy/administer
  - if i pgp an email to one of my favorite security geniuses, it adds a week to RTT
    - **if he reads it at all**
    - **and that's only if i manage to have a version-compatible key**
    - **will admit i do same for ppt and doc files. but why punish ascii?**
    - **why don't all client mail handlers support transparent authentication/encryption?**
  - even for an adolescent industry this component fails the smell test
    - **and we're a little beyond that now anyway**
    - **teenage scars notwithstanding, we are capable of cooperation**
  - we shouldn't indulge this splintering unless it offers benefit
    - **if anyone defends the genetic diversity of the pgp landscape i will personally flog them**

</rant>

problem of the Internet

qos: mechanisms to differentiate performance based on application or network-operator requirements

- or provide predictable or guaranteed performance to applications, sessions, or traffic aggregates
- innovations emerged in several areas
  - packet scheduling, admission control, traffic shaping
- successful in constrained scenarios
  - VOIP
  - empirical load-based capacity planning
- wrt interdomain, it went as far as it could go technically without economic (network market) support
  - a few years in the lab can often save a few hours talking to a provider
  - economic technology just not there
- also huge sociocultural resistance to paying more
  - users think the Internet should just work
  - they've seen it happen before
  - tools to separate and service differentiate topologies now emerging in protocol specs
  - but still insufficient market support

problem of the Internet

compromise of the e2e principle

- "do not replicate in lower layers what can be handled by higher layers"
  - has taken a beating this decade (and it's still early)
- in its place, a web of contracts to control what people are allowed to do
- saddest part: we had a different solution (IPv6), but too little too late
  - NATs, firewalls demanded viscerally by the market
  - the same brilliant community ultimately brought you both in (some defn of) parallel
- we have the Internet [un]layering we deserve
  - it's a mess to be sure, e.g, IPSEC through NATs/firewalls
  - "sometimes the price of freedom is what freedom brings" -- eric schlosser's **reefer madness**
  - we have failed you (sysadmins) by engineering our way toward the unsupportable
  - maybe if more sys and net admins had been in some of those IETF WG meetings....

need to be realistic about where to go from here

- IPv6 will be hard-pressed to revive e2e legitimacy (tho it has believers)
- don't think we're getting the e2e architectural assumption back
  - need to think outside that box from now on
- right solution at time t might not actually be the right solution at t+1
- not too late for admins to get [back] in on the fun

problem of the Internet

dumb network

- dumb as in mute -- as in it can't talk to us about its internal state
  - can't tell us how much bandwidth it has
  - can't tell us why it changed its route
  - can't change a route because we want it to
  - can't tell us if it's being attacked
  - for something built for communication, it's pretty disappointingly uncommunicative
  - makes it hard to manage and provision/engineer its growth
  - no wonder we engineer blind so much of the time
- ok but there's dumb and there's idiotic
  - a routing architecture that requires humans in various NOCs to tweak link weights for good performance?
  - that can't have been "plan A".
- need for greater Internet transparency
  - grenville's nice talk on the topic http://caia.swin.edu.au/talks/031002A/031002A.pdf

largely goes back to measurement. the energy we've invested in measurement of this infrastructure is far less than has been invested in any other aspect of this infrastructure, and now we're wondering why we're having a hard time getting a handle on it.

[kc: not that i'm bitter :)]

problem of the Internet

robust scalability of routing system

- closely related to configuration management
- primary factors in routing evolution:
  - relative cost-performance of communication, computation, and human brains
  - tradeoff between fast convergence and stability for current IGPs
    - **timers limit effect of external instability at expense of increased convergence time**
    - **hard to get data to do real studies/analysis to discern real from artificially imposed instability**
    - **better damping algorithms remain elusive**
  - researchers & sysadmins can help optimize navigation of trends
    - **less hope of changing them**
- worse news: we really don't understand the design space
  - routing architecture stagnates (unless you count the hacks)
  - no way to judge success or failure of proposed architecture
  - or verify operational integrity
  - any change sufficiently ambitious to address problems is also sufficiently ominous to scare any vested interest whose support is required
- in meantime, problems with current routing have not even begun
  - BGP has no mechanism to route around saturated chunks of core
    - **core Internet chunks operate for weeks/months at/near capacity**
    - **too much manual tweaking going on to justify an assumption that hell won't break loose at some point**
  - proposed overloading of BGP infrastructure to distribute "non-routing" information"
    - **auto-discovery mechanisms for Layer 3 VPNs [BGPVPN]**
    - **not particularly comforting that we're adding even more responsibility to a system we don't really understand**

robust scalability of routing system (2)

difficult to get consensus even (esp) among routing experts:

(akamai researcher bruce maggs, routing wksp www.net.informatik.tu-muenchen.de/wired/)

- 1. Where (if anywhere) is the congestion in the Internet?
- 2. How much capacity does the Internet have, and how fast is it growing?
- 3. How much traffic does the Internet core carry and what does it look like?
- 4. How fast is network traffic growing?
- 5. What will traffic patterns look like five years from now?
- 6. Can we scale the network to support the demands of users five years from now?
- 7. How much does/will it cost to increase network capacity?
- 8. Will stub networks soon be employing sophisticated traffic engineering mechanisms on their own, e.g, those based on multihoming and overlay routing? What impact might these techniques have?
- 9. What fraction of traffic are CDNs carrying? What effects derive from DNS tricks to route traffic?

robust scalability of routing system (3)

this actually gets pretty ominous

- many believe the routing system will find a state of non-convergence that so disruptive as to bring down large portions of the Internet
  - we talk about malice, but more frightening truth is that we aren't sure a typo couldn't do this
  - we can't even trace back DOS attacks
  - debugging remains black art
  - routing protocols interact with eachother in 'interesting' (nonunderstood, sometimes nondeterministic) ways
  - intelligent routing throws a wrench into the melting pot
  - scalability and robustness require even more 'damage control' complexity
- perfect 'normal accident' suggested possible by bgp expert tim griffin
  - where no single ISP will be able to identify and debug the problem
  - where it will take days to fix and cost the world economy billions of dollars
  - where the press will learn that the internet engineering community had known about this lurking problem all along....
  - for front-row seat at melange of fingerpointing, keep an eye on Internet routing system
  - promise you (sysadmins/netadmins) won't be left out
- in the meantime you have enviable job security (you're welcome)
  - and an excruciating if not impossible job (oops, sorry)
  - 3600 RFCs later and your job gets harder rather than easier each day
    - **what is wrong with this picture**
    - **rfc used to stand for something...**

robust scalability of routing system (4)

what is needed (courtesy tim griffin, phil karn)

- defined routing policy languages guaranteed to be globally sane
  - no matter the what local policies are defined
  - BGP speakers must be forced to use them (i.e, standardize MUST)
- give user control of packets to/from his/her own IP address
  - rather than clumsy, brittle firewalls not understood by the ISP's phone support anyway
  - open standard for the secure remote control of a generic packet-filtering firewall

research questions

- is it possible to design such languages and protocols?
- how can we find the right balance between local policy expressiveness and global sanity?
- what exactly do we mean by "autonomy" of routing policy?
- do we need additional protocols to enforce global sanity conditions?
- how can we enforce compliance of policy language usage?

robust scalability of routing system (4)

what sysadmins can do now

- most important: don't assume we have this under control
- read through some of geoff huston's work -- great introduction
  - http://bgp.potaroo.net/
- get involved w the IETF
  - ask if new routing products/services make things better or worse for the commons
  - ask why the underlying architecture doesn't obviate the need for hacks
  - management willing (i am aware that many of you are already doing N FTE's of work)
  - strategic involvement (ietf can be imprudent use of your time. get experienced mentor)
- use routeviews.org to look at the routing system
  - it was built for you
- document your own topology internally
  - at more than one layer
  - often
- work with researchers: data analysis, visualization
  - lots of them are looking for good problems
  - you definitely have some
  - consider it a chance to save them a few years of irrelevant research

problem of the Internet

normal accidents (not just a recommended book anymore)

- dave plonka's talk from yesterday
  - several other examples of hard coded IP addresses wreaking havoc
- DNS: chock full of inability & need to evaluate macroscopic performance
  - caida rfc1918 paper, effect of anycast traffic, etc
- common now to deploy a half a million homogenous Internet hosts
  - low price point
  - dramatic change in Internet OS landscape
  - what happens when each bic lighter has an IP address?
- market pressure forestalls adequate testing
  - not that we even know what that means
  - testing for tomorrow's Internet: an intractable task
- lack of body for specification and conformance with RFC-defined standards by designers/manufacturers/vendors
  - http://www.ietf.org/internet-drafts/draft-loughney-what-standards-00.txt
- no underwriters laboratory (ul.com) for things that talk to the Internet
  - includes fighting back when needed measurement functionality is unsupported
  - who would take this on?
  - if this doesn't happen on its own, will some www.dhs.gov spirit force it?

normal accidents (2)

why this problem is important

- the industry will recover
  - including from its post traumatic stress disorder
- and these normal accidents will increase in number and ramifications (cost)
- stakes grow monotonically

what can be done to help

- need labs that mimic your your infrastructure
  - smaller bandwidth ok
  - work with your vendor, e.g, cisco provides labs for software upgrades
- collaborative efforts among operational and research communities
  - isc
  - pch
  - routeviews
  - ripe
  - nlnet
  - caida

problem of the Internet

lack of quanitified threats to infrastructure:
no rigorous study exists of root causes of Internet performance problems/outages

anecdotal survey (courtesy sean donelan nanog post):

1. network engineers (what's this command do?)
2. power failures (what's this switch do?)
3. cable cuts (backhoes, enough said)
4. hardware failures (what's that smell?)
5. congestion (more bandwidth! Captain, I'm giving you all she's got!)
6. attacks (malicious, you know who you are)
7. software bugs (your call is very important to us....)

"knowing what we don't know" offers little comfort

this one mostly out of scope for sysadmins and researchers

- fcc will have to get involved

problem of the Internet

intellectual property & digital rights

- obviously a layer umpteen issue but it's too hot not to touch
- lessig covers this topic quite well
  - if you haven't read his books yet, do so over the holidays
- internet not a dichotomy between commercialization & open standards
- trichotomy, with 3rd piece: regulation, property rights and protection
  - sysadmins will be in position of reading subpeonas from RIAA (some of you already have)
- instantaneous sharing: best and worst thing about the Internet
  - biggest threat is the "downside of the upside"

what can be done to help

- not blocked on technology (or computer science. or sysadmins)
- blocked on our lack of social consensus how to incorporate the reality of digital ubicopyright into our model for appropriate human interaction
  - we should keep "starving artist" two words
  - not just about artists. as sterling would point out, we haven't sorted out the impact
  - e.g, neglected side effect: increases effective value of in-person contact
    - **live concerts/interactions now become vastly more valuable than the next step down**
    - **even for people not using their PDAs to look up band facts or check mail**
- so talk/write to/for legislators. tell them what you know
  - else they decide without us

problem of the Internet

governance *(see vixie's talk yesterday, www.caida.org/outreach/presentations/governance2003/)*

- shared resources need global administration
  - universally agreed allocation of addresses, ASes, domain names, protocol numbers
  - pending bill manning's multiple-NATed Internet
    - **proposed backbones use RFC1918 space, all customers use rest of IPv4 space as private space**
- considered harmful: operating heavy machinery under the intoxicating influence of mindbending revenue potential
  - like .com
  - heaven forfend the dns root system
  - sitefinder and countermeasures fall into cyberwarlordism category
  - can we (socially) increase the set of parties (besides shareholder) that an Internet company considers a constituency
  - because that's what it will take
  - but at least sitefinder put the 'steward' vs 'owner' issue on our kitchen table where it belongs!

what can be done to help

- may sound a little old by now but: participate!
- icann mail lists, www.icannwatch.org, www.arin.net for address policy
- join local isoc, go to arin and icann meetings (all open)
- the policy process has been taken away from us less than we think
  - still more than we wish. but it's no excuse to withdraw
  - no proposed option has been better

problem of the Internet

growth in traffic and user expectations

- bandwidth budgets frozen but application designer creativity not
  - VOIP, IM, P2P, streaming
- data center consolidation
  - move servers away from users, often puts traffic across WAN
- more users!

user expectations (and about those users)

- users are abstracting way faster than we can
  - they know all this stuff should work any day now
- they are ready to configure things we don't even know how to describe much less support
  - e.g, please give me this much quality of service for this long
- they want 'tennis racket' transparency
- right now we have the recession as an excuse
  - no capital for infrastructure expansion
- as recession subsides, user expectations will increase

am not going to offer solutions to these problems because i rather like having them.
why should anyone miss out on the Internet?

problem of the Internet

interprovider (& vendor/business) coordination

- companies avoid publicizing a vulnerability of their own infrastructure
  - silence renders overall system more vulnerable

what can be done to help

- need measurement repositories for data to support debugging
- make friends with researcher, or provider
  - refine requirements/approaches, and cost models
- don't underestimate the value of cross-fertilizing your brains
- avoid govt accusations of antitrust activity by including them
  - DHS will make this 'easier'
    - **</cough>**

problem of the Internet

time management/prioritization of tasks

- strongly interrupt-driven field
- sometimes need allocated time just to think
- plan strategically rather than tactically
- unfortunately this problem pervades all levels of Internet design
  - i hear other fields have it too
- lisa gives tutorials on this; you guys likely know more about it than i
- 'overspecialization' was also mentioned
  - "solaris team doesn't upgrade solaris NIC, that's networking's job"
  - undoubtedly a large company phenomenon
  - inherent tension between automation and job security

enough problems -- let's talk doomsday

- "the Internet is dying" -- Karl Auerbach provocative article
  - http://www.circleid.com/members/profile_view_ind.php?id=29
  - 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
    - **i know some of you wouldn't consider this a problem**
  - in its place a much more controlled approved set of communications
  - lesson 1: don't run tcpdump if you don't want to get depressed -- most of it is garbage
    - **last i checked most TV was garbage too. is it losing users? or do we get smart technologies to help us use it**
- "digital imprimateur" -- john walker
  - http://www.fourmilab.ch/documents/digital-imprimatur/
  - "how big brother and big media can put the Internet genie back in the bottle"
  - rich 'optimistic pessimism'
- larry lessig's code of laws and future of ideas (from 4 slides back)
  - by leaving policy to the policy folks your future derives directly from their clue level
    - **we need to own up to that**
  - most optimistic pessimist lawyer on my bookshelf
- bruce sterling keynote at NSF workshop feb 2002
  - http://www.cra.org/Activities/grand.challenges/sterling.html
  - all SF writers are optimistic pessimist so that's not his accomplishment
  - his writing is...exceptional [read quotes]
    - **ubicomp and ultrawideband and machines-building-machines are his messiahs**
  - he doesn't hold back against the computer industry. cute.

obligatory quote from castells

writing a trilogy, quote from first volume below

- not the easiest reading, but we all should anyway

//it is the beginning of a new existence, and indeed
the beginning of a new age, the information age,
marked by the autonomy of culture vis-a-vis
the material bases of our existence.
but this is not necessarily an exhilirating moment.
because, alone at last in our human world, we shall
have to look at ourselves in the mirror of historical reality.
and we may not like the vision.//

-- manuel castells, rise of the network society, vol1, p. 478

so what now

patience, persistence, and perspective

- bruce sterling described these in 1994 as most the important virtues for the computer industry to embrace.
- 'normally somewhat dull virtues'
- 'the most difficult to manifest from within a revolution'
  - if you think the revolution is over you aren't paying attention
- especially when you're all shouldering generations of neglected architectural responsbilities
  - i did apologize for that already...
- politically minded disposition never hurts

involvement with policy and research communities

- help educate at least one

so what now (2)

awareness of your role

- sysadmins are key (and unsafely ignored) channel between R&D community and real world
  - shrapnel-closest to real problems
  - hard-earned intuition and insight that we don't have
- consider participating in the policy process that will render you cogs of big brother if you don't
  - e.g, IETF wiretapping issue in februrary 2002
- you are dangerously underrepresented in these communities
  - this means NANOG, ARIN, NANOG, IETF, research workshops and conferences
  - all these people are playing with layers that you have to manage
  - and they are people sometimes more loyal to interests other than the Internet
- educate your management
- check in with the research community every so often
  - most of us hit the bar every so often so you probably don't have to walk that far
  - and we need good relevant problems
- also, reminder: continually watch your network
  - still be confused but on a higher level

so what now (3)

for USENIX/LISA

- expand your role as an operations research forum
  - more peer-reviewed research
  - court the traditional network research community
  - court the operational networking community
  - joint workshops on configuration management problems
  - so ops folks can give feedback on how research ideas interface with current operational reality
    - **save them a few years**

conclusion

sorry for the empowerment speak

- wouldn't get so lofty if you didn't hold the ring of power, frodo
  (at least i hope you have it because i've asked around & nobody else thinks they have it)

- Internet has done a phenomenal job at dramatically reducing the space in between people who want to communicate
- next job is to reduce the space between the people who want to communicate and the Internet
- that's an increasing proportion of your job today
  - it will continue to increase in the future
  - if part of that means making parts of your job unnecessary, we promise we'll make more messy technology
  - making parts of your job unnecessary should be your overriding professional concern for this decade

- if sysadmin job security ever becomes a top problem of the Internet, i promise i'll give another invited talk on what you should do next

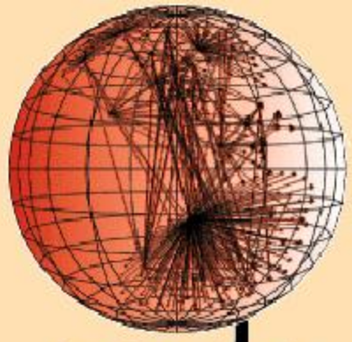

// disorder increases with time
because we measure time in the
direction in which disorder increases. //
-- stephen hawking

kc@caida.org
nov 2003

appendix: trilogy of action for scientists

for 'seekers of the larger view'

- draw together pieces of science and technology to create a system
  - **whether that system is xerography, telegraphy or steam navigation.**
- find the economic feasibility for a new technology
  - **by virtue of a wide grasp of the worlds of man and matter**
- reach harmony through intuition
  - **by meditating on deep knowledge of the field so as to arrive at a new result**
- build a model
  - **a simplified representation of the problem, subject to experimental analysis**
- serve as a science-technologist generalist
  - **who, many times/year, extracts the missing point out of a complicated situation**
- make decisions or help others make decisions
  - **by imaginative interaction w/alternatives calculated as consequent on those decisions**

-- john archibald wheeler