Cooperative Measurement and Modeling of Open Networked Systems (COMMONS) proposal

cooperative association for internet data analysis

26 october 2006

outline of talk

problem(s) statement: Internet in crisis!

and various reactions from stakeholders

historical context (how we got there)

what have we learned from studying the Internet and how can we apply it?

proposed solution(s): COMMONS

what we (all) can do to help

urgent problems

(1) commercial sector reports financial crises that threaten first amendment

(2) emerging community networks lack resources and experience to make informed provisioning decisions, continually threatened by incumbent-driven legislation

(3) no incentive or funding even for public sector to provide access to data on operational infrastructure. so researchers can’t help

more (and less) urgent problems

(4) running out of IPv4 addresses, intended solution not meeting requirements
(5) routing system hitting fundamental limit
(6) best traffic engineering exacerbates it
(7) Internet killing the phone business
(8) regulatory framework a mess
(9) dismal field of network science stunted

(10) we can’t even have an empirically grounded conversation about it

May 24, 2006: Deep IMPACT: Justifyin

About SANS | Contact SANS | Wha

Electronic Newsletters | WhatWorks

The Twenty Most Critical Interne

Version 6.01 Nov

Questions /

To link to the Top 20 List, use

AccessIT GROUP

-----Jump To Index of Top 20 Vulnerabilities -----

Introduction

The SANS Top 20 Internet Security Vulnerabilities

Four years ago, the SANS Institute and the Nation

Cente

Intern

expan

their

servic

these

This S

additi

Applic

the ev

the pr

listed

patche

patch

We ha

Hence, the Top-20 2005 is a co

Home >> China

China adds top-level domain names

China's Ministry of Information In

domain name system in accorda

Regulations.

After the adjustment, ".MIL" will b

"CN".

A new Internet domain name sys

Under the new system, besides

"NET" are temporarily se

servers under the management

Numbers (ICANN) of the United

The In

securit

can he

and m

Sear

Resp

Opinio

Shouldr

High

The ope

that cou

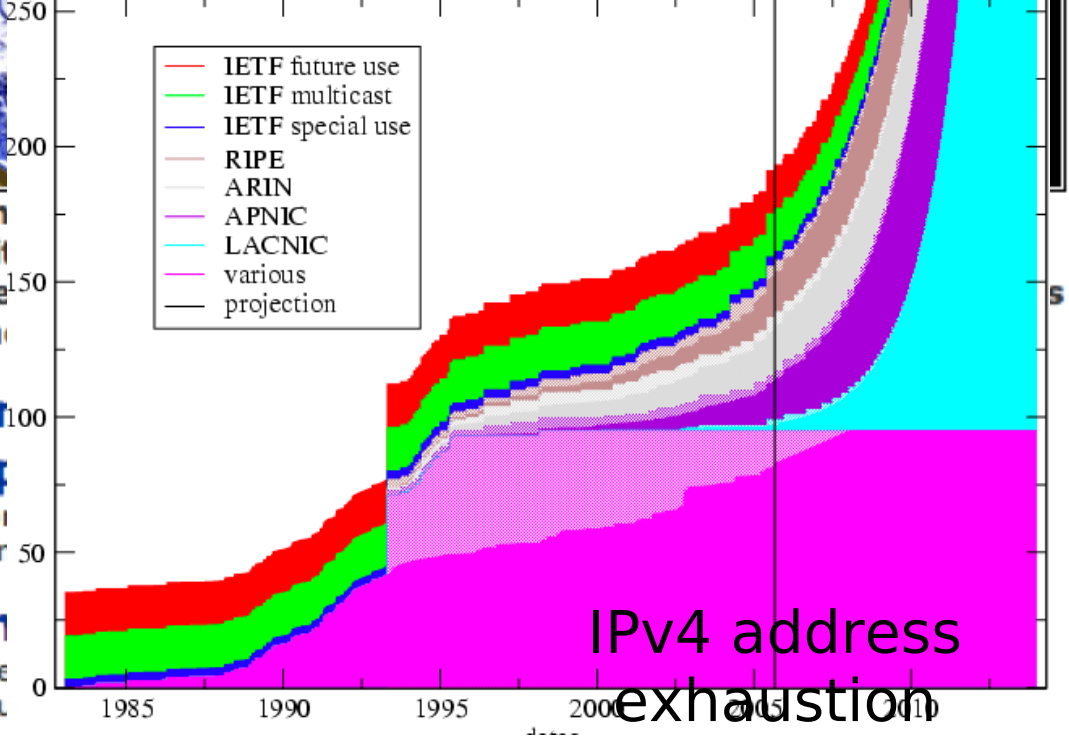

IPv4 address exhaustion

The Dark Side of the Search Engine Business

Paid search is a booming business for Google, Yahoo and Microsoft, but there's a major downside for users. A new study by McAfee's SiteAdvisor finds sponsored search results contain two to four times as many dangerous

How the internet killed the phone business

Almost-free internet phone calls herald the slow death of traditional telephony

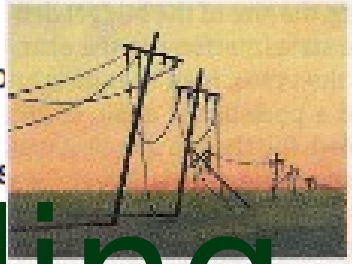

THE term "disruptive technology" is popular, but is widely misused. It refers not simply to a clever new technology, but to one that undermines an existing technology—and which therefore makes life very difficult for the many businesses which depend on the existing way of doing things. Twenty years ago, the personal computer was a classic example. It swept aside the older mainframe-based style of computing, and eventually brought IBM, one of the world's mightiest firms at the time, to its knees. This week has been a coming-out party of sorts for another disruptive technology, "voice over

market, as the marginal price of making phone calls heads inexorably downwards.

VOIP makes possible more than just lower prices, however. It also means that, provided you have a broadband connection, you can choose from a number of providers of VOIP telephony and related add-on services, such as voicemail, conference calling or video. Many providers allow a VOIP account to be associated with a traditional telephone number or with multiple numbers. So you can associate a San Francisco number, a New York number and a London number with your computer or VOIP phone—and then be reached via a local call by anyone in any of those cities.

Furthermore, your phone (or computer) will ring wherever you are in the world, as soon as it is plugged into the interne

falling bits of sky

International Telecommunication Union

Building Partnerships for Progress

All match

More

Home : Office of the Secretary General : SPU

Next Generation Networks

Home: OECD > OECD ICCP Workshop: "The Future of the Internet", Paris, 8 March 2006

OECD ICCP Workshop: "The Future of the Internet", Paris, 8 March 2006

Send Print

IPv6

From Wikipedia, the free encyclopedia

Internet Protocol version 6 (**IPv6**) is a network layer standard used by electronic devices to exchange data across a packet-switched internetwork. It follows IPv4 as the second version of the Internet Protocol to be formally adopted for general use.

IPv6 is intended to provide more addresses for networked devices, allowing, for example, each cell phone and mobile electronic device to have its own address. IPv4 supports 4.3×10^9 (4.3 billion) addresses, which is inadequate to give one (or more if they possess more than one device) to every living person. IPv6 supports 3.4×10^{38} addresses, or 5×10^{28}(50 octillion) for each of the roughly 6.5 billion people alive today.

Invented by Steve Deering and Craig Mudge at Xerox PARC, IPv6 was adopted by the Internet Engineering Task Force in 1994, when it was called "IP Next Generation" (IPng). (Incidentally, IPv5 was not a successor to IPv4, but an experimental flow-oriented streaming protocol intended to support video and audio.)

THE INTERNET.COM

Fight for Internet Freedom

the coalition f.a.q. press

THE LATEST....

Moby Speaks Out on Internet Freedom

At a press event in Washington today, Grammy-nominated musician Moby (along with

RKWORLD RESEARCH CENTER: Convergence / VoIP

IP PBX SIP VoIP Services Vendor Solutions

NetworkWorld.com > Convergence / VoIP >

What IMS promises enterprises and carriers

Internet Protocol Multimedia Subsystem called key to converged, expanded services.

By Stephen Lawson, IDG News Service, 09/26/05

where we are
[Dr. JCarl
patibility and
"where we
plans." "I the
gure out how b

sundry "solutions"

public sector starting inquiry

knocking at CAIDA's door with questions about the Internet:

DHS, NCS, NIST, DOE, FCC, FTC, NSA, GAO, NSF (and that's just the US government..)

entire muni and community wireless networking movement...

Community Networks Inside the US

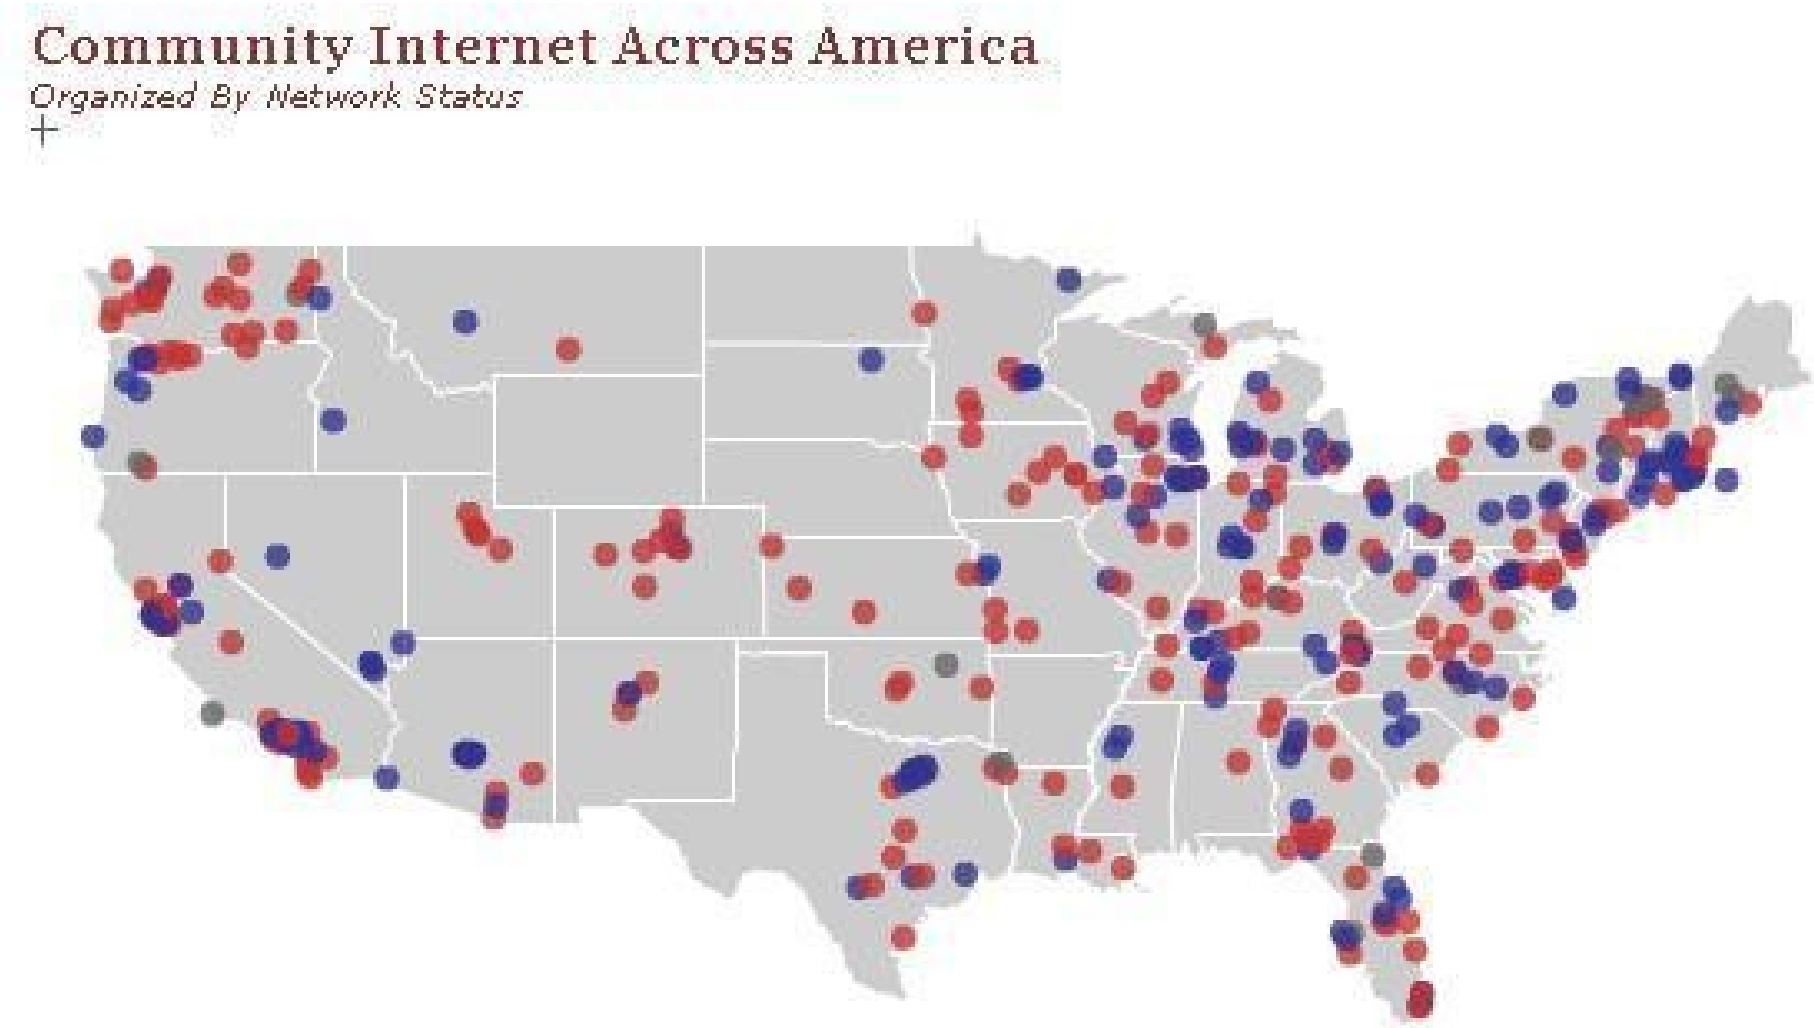

Graphic Credit: Free Press

The Future of the Internet

In a decade, the Net will dig deeper into our lives.

April 10, 2006 Issue

Credit: Dave Cutler

http://www.redherring.com

"While the business case for the carriers may be disappearing, a host of new business and investment opportunities is being created with far greater economic wealth creation," Mr. Arnaud writes in his blog. "Our biggest concern is that governments will be distracted by the complaints of the old industry such as carriers and penalize the new economy industries of the Internet."

"We don't presently have a roadmap of where we are trying to go with the Internet," says MIT's Mr. Clark. Instead of worrying about backward compatibility and migration issues, the focus has shifted to "where we would like to be in 10 to 15 years," he explains. "If the story is compelling enough, people will figure out how to get there."

National Science Foundation
DIRECTORATE FOR
Computer & Information Science & Engineering (CISE)

CISE Home | CISE Funding | CISE Awards | CISE Discoveries | CISE New

mputer & Information
iences & Engineering

I Home
I Announcement
quently Asked Questions
kshop Reports
sentations
r Information

The GENI Initiative

The Directorate for Computer and Information Science and Engineering (CISE) is planning an
Environment for Networking Innovations or GENI to explore new networking capabilities that
stimulate innovation and economic growth. The GENI Initiative responds to an urgent and imp
Century to advance significantly the capabilities provided by networking and distributed syste

The GENI Initiative envisions the creation of new networking and distributed system architectu

- Build in security and robustness;
- Enable the vision of pervasive computing and bridge the gap between the physical
  mobile, wireless and sensor networks;
- Enable control and management of other critical infrastructures;
- Include ease of operation and usability; and
- Enable new classes of societal-level services and applications.

The GENI Initiative includes:

- A research program; and
- A global experimental facility designed to explore new architectures at scale.

CISE is encouraging a broad community effort that engages other agencies, other countries, a

THE GENI RESEARCH PROGRAM

(US) NSF's hand

e.g. NSF's GENI initiative

US NSF responding to network research community frustration

- difficulty with technology transfer, not to mention science
- persistent problems leaking into unready world

attempt to redesign components 'in the light'

what did we learn from measuring this one?

NAS report on ‘network science’

1) networks are everywhere and thus important

2) we don't yet have any predictive power over complex networks

3) funding situation backwards: domain-specific (splintered) rather than fundamental

NAS report on ‘network science’

identifies as top three challenges:

1) characterization of dynamics and information flow in networked systems

2) modeling, analysis, and acquisition of experimental data for extremely large networks

3) rigorous tools for the design and synthesis of robust, large-scale networks

http://fermat.nap.edu/books/0309100267/html

NAS netsci recommendations

1) federal govt initiate focused R&D program to close the gap between currently available knowledge required to sustain the complex global networks on which the well-being of the US has come to depend
2) Army should invest $10M/yr in ways different from other agencies
3) basic research program on interaction between information networks and social networks that use them

http://fermat.nap.edu/books/0309100267/html

context

critical infrastructure

what is it? how does it get that way?

what are common characteristics?

is the Internet one? or will it be soon?

what are the implications for public and private sectors?

underlying goals: innovation, economic strength, democracy, freedom, health, science, arts, society.

it really is about living in a better world...

16 operational internet problems

- security
- authentication
- 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
- stewardship vs governance
- intellectual property and digital rights
- interdomain qos/emergency services
- inter-provider vendor/business coordination

persistently unsolved problems for 10+ years

(see presentations at www.caida.org)

why we're not making progress

- if providers are broke, they can't invest in long-term health of infrastructure.
- so add to list of problems: **sustainability**
- top unsolved problems in internet operations and engineering are rooted in **economics, ownership, and trust (EOT).**

does not mean there aren't useful technical problems to study. but there will be no technical solutions to these problems that don't solve the EOT issues.

historical context

1966: Larry Roberts, “Towards a Cooperative Network of Time-Shared Computers” (first ARPANET plan)

(*we are still using the same stuff*)

1969: ARPANET commissioned by DoD for research

1977: Kleinrock’s paper “Hierarchical Routing for large networks; performance evaluation and optimization”

(we are still using the same stuff)

1980: ARPANET grinds to complete halt due to (statusmsg) virus

1986: NSFNET backbone, 56Kbps. NSF-funded regionals.

IETF, IRTF. MX records (NAT for mail)

1991: CIX, NSFNET upgrades to T3, allows .com. web. PGP.

1995: under pressure from USG, NSF transition of backbone to competitive market. no consideration of economics or security. kc proposes caida.org. universities build v “Internet2”...

2005: *The Economist*’s cover story: “*How the Internet killed the phone business*” (September)

what have we done?

we replaced a critical infrastructure with something not designed to be critical infrastructure

historical context explains it but does not address incongruities

and this decade, free markets go up against free speech

what have we learned?

most important thing we've learn so far: society has decided IP is like water.

> *"our best success was not computing, but hooking people together" --david clark, 1992 ietfplenary*

strong implications for an industry structuring itself to sell wine. but that's what the data shows.

when you want to move water, you care about 4 things: safe, scalable, sustainable, stewardship.

the 4 S’s

- safety: is the data toxic upon arrival?
- scalable: can we route/name/address earth’s needs?
- sustainable: is it economically viable?
- stewardship: will the provisioning and legal frameworks we choose leave our children -- and democracies -- better or worse off?

none are purely technical, but all require technical understanding to get right.
and they’re all connected.

how have we done?

- how safe is the Internet?
  - data doesn’t look good
- how scalable is the Internet?
  - data doesn’t look good
- how sustainable is the Internet?
  - data doesn’t look good
- how did we do on stewardship?
  - data doesn’t look good

not that we haven’t been trying

e.g., all caida projects are on the 4 S’s:

1) safety: security, DNS, PREDICT, telescope
2) scalability: routing and topology research
3) sustainability: EOT, DNS, COMMONS
4) stewardship: address consumption, trends, all measurement & data activities

measurable progress on real Internet eludes us

failure (to measure progress) on 4S's poses risks to economies & democracies:

- that we won't learn from our own history, won't admit we don't understand the economics, and thus must set policy based on unvalidated assumptions
- that we will design another architecture with no actual plan for economic sustainability (much less incenting further innovation in a competitive market!)
- that other forces will "code" innovation into the architecture (free markets vs free speech)

there is good news

- we made something so great, everyone wants it.
- in fact many of us want it more than once! (um..)
- the current industry is a historical artifact of technical and (science & regulatory) policy ‘innovations’ in the 60s, 70s, 80s, 90s, and 00s
- people are starting to study interplay, but they’re undercapitalized
- in the meantime, it became global critical infrastructure. oops.

network economics: dismal science(s)

known: economics of current architecture need study

- have never been a priority.
- conversations for last 15 years have been private
- enlightened policy impossible

our misunderstanding the economic architecture threatens an architecture we hold much more dear..

time for the academic community to step forward!

proposal

Cooperative Measurement and Modeling of Open Networked Systems (COMMONS)

(1) offer cooperative backbone in exchange for mutual, privacy-respecting, community-defined transparency across network

(2) experiment with different architectures: not just technical, but economic, ownership, trust

(3) use strengths of Internet to overcome its weaknesses

NLR-peered network

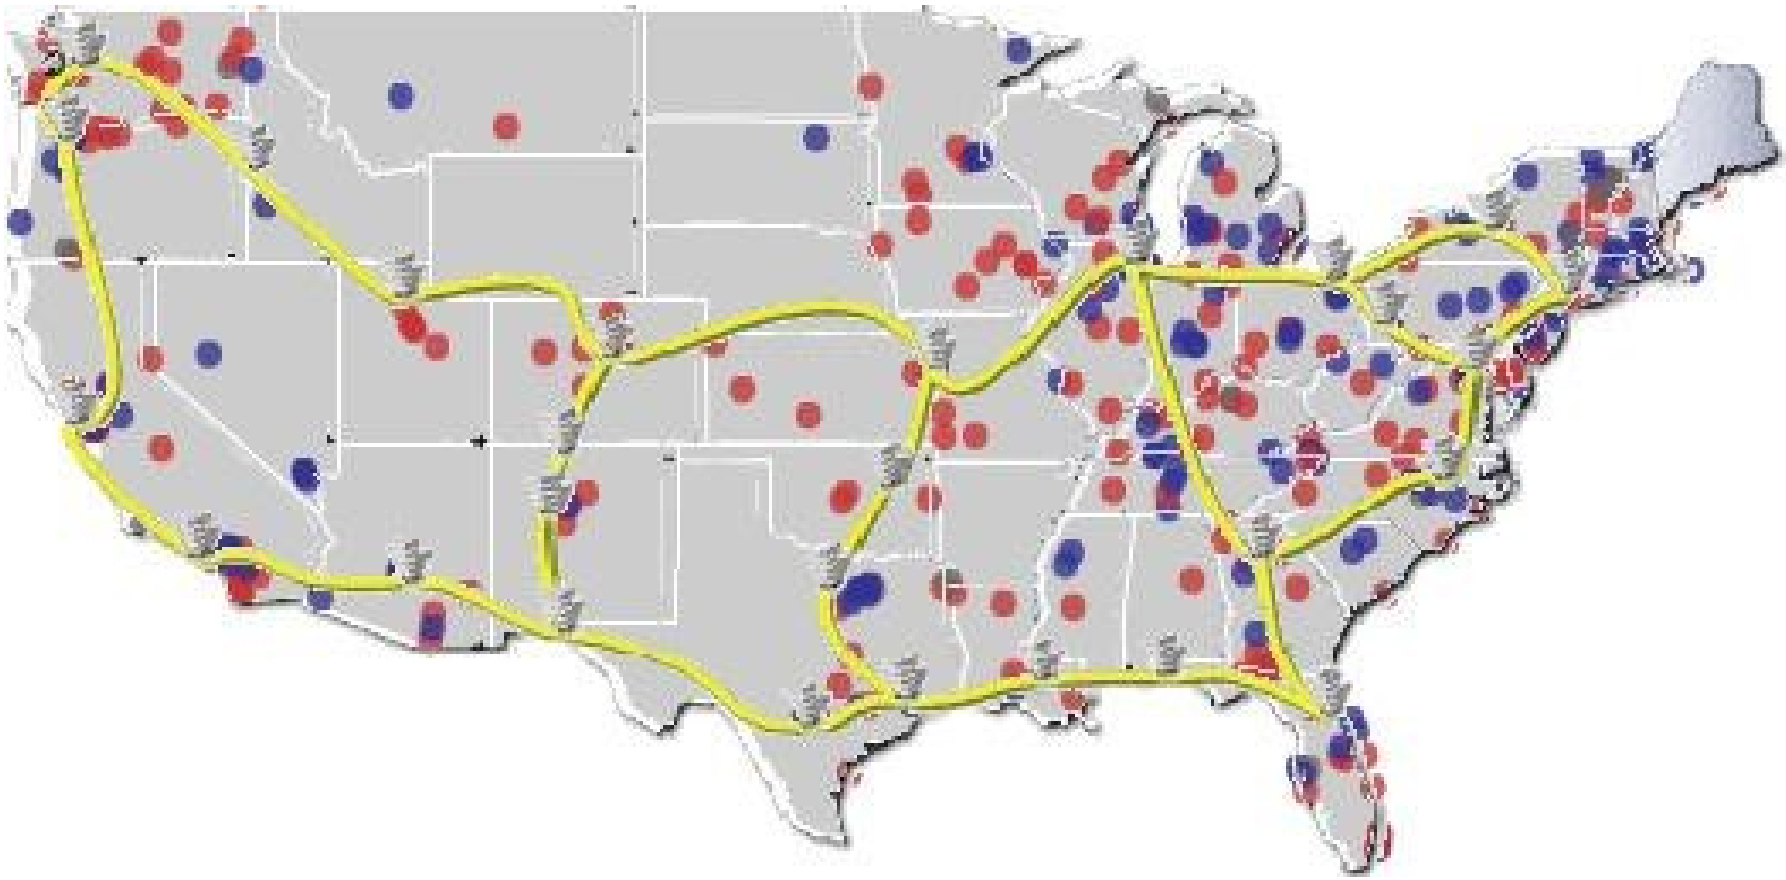

Graphic Credit: Free Press/NLR/Carl Malamud

measurable progress

solving acute problems (slide 3)

(1) alleviate commercial sector of impossibly low margin business...and secure first amendment

(2) provide emerging community networks with level playing field, and critical mass of expertise from which to draw

(3) gives science a chance; creates a resource for network science for the public good

measurable progress

what about the other problems?

(4) running out of IPv4 addresses, intended solution not meeting requirements
(5) routing system hitting fundamental limit
(6) best traffic engineering exacerbates it
(7) Internet killing the phone business
(8) regulatory framework a mess
(9) dismal field of network science stunted

** **but at least we can start an**

empirically grounded conversation **

measurable progress

solves longer term problems too

(1) creates opportunities for sound measurement and analysis – key telecom policy that serves public good

(2) helps achieve universal affordable service that free market has failed to deliver

(3) facilitates solutions that push control (and economics) as far to edge as possible

(4) foster new generation of innovation in service, applications, hardware & software

broader impact

multi-stakeholder approach

(1) address immediate policy concerns

(2) enables revolutionary educational experimentation

(3) supports public sector networking experiments, e.g., disaster response, community watch, civic debates

(4) establishes a path to a science of cyberinfrastructure

potential partners

- NSF/GENI
- Internet2
- QUILT
- NLR
- RONs
- Educause, NATOA, & other coalitions
- state networks
- municipalities/community Wifi implementors
- CRACIN & similar organizations

http://www.caida.org/projects/commons/