Internet as emerging critical infrastructure: what needs to be measured?

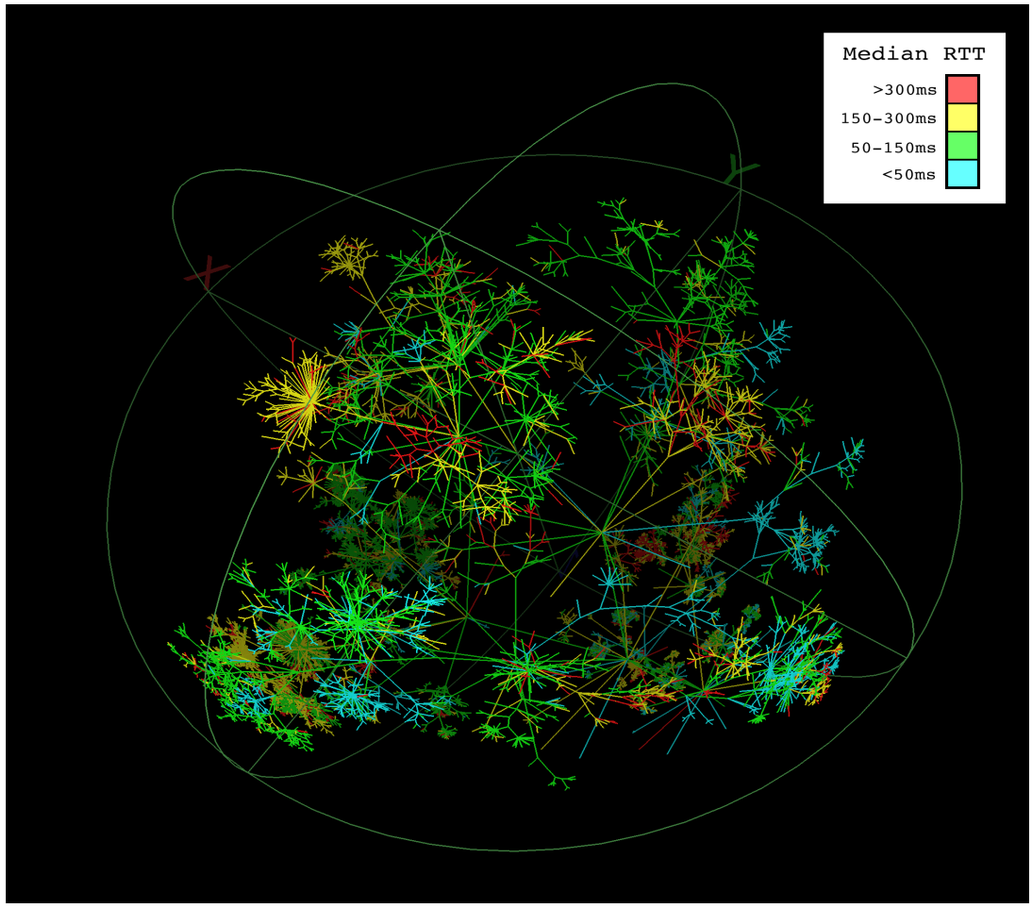

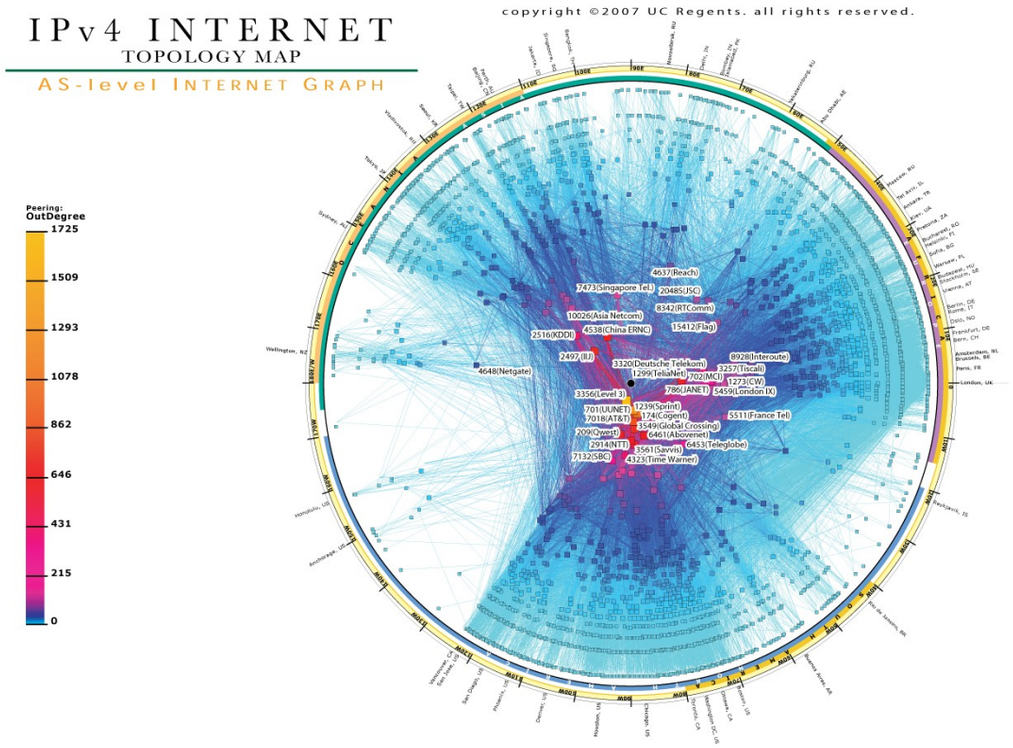

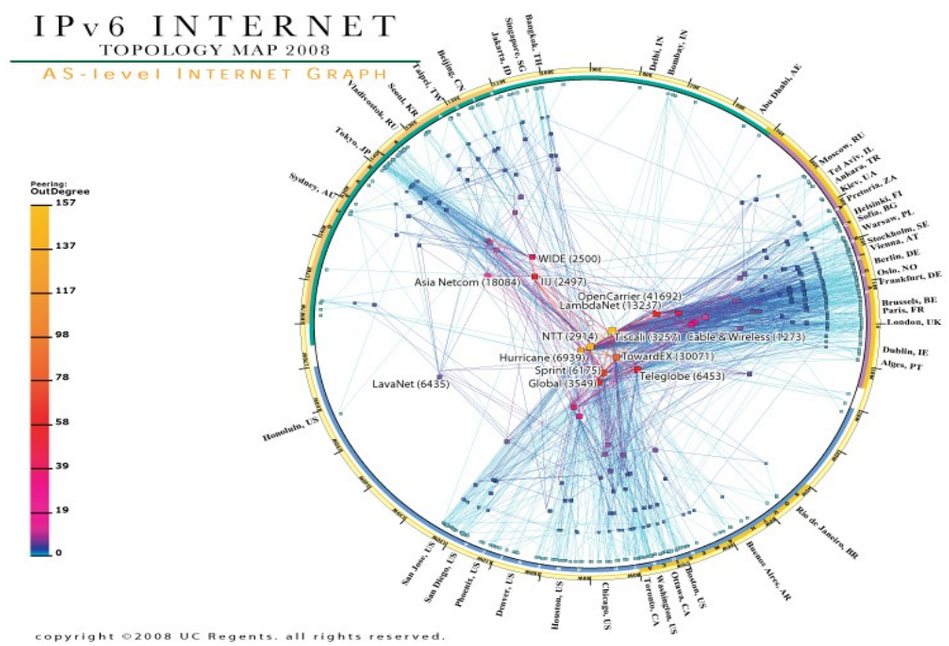

kc claffy

kc@caida.org

presentation at JCCC08
12Nov 2008

recipe for disaster

(aka “you are here”)

- We now critically depend on the Internet for our professional, personal, and political lives.
- But what do we know about it? e.g, what keeps the system stable or drives it to instability? How can we protect it from operational threats?
- Researchers and policymakers currently analyze a trillion dollar industry in the dark.
- Few data points available suggest a dangerous picture.

How did we get here?

- Telephone system: 140+ years of history, including regulated data collection requirements (and profits). and a precisely defined system.
- Data networks: 40 years old, ad hoc/hack, tossed to private sector before mature, with no government support for research or metrics (or profit), ill-defined system.
- Current academic projects either lack sustainability or ability to dedicate resources
- War: in U.S. the best motivation so far for investing in understanding critical infrastructure

CAIDA: background & history

Since 1997: narrowing the gap between Internet operations and science in face of global privatization

Largely US taxpayer funded (nsf, dhs), plus sponsors

Seek, analyze, communicate relevant features of best available data on the Internet

Use this data to prepare for the future

Recent expansion of research agenda into policy and economics

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 Center (NIPC) at the FBI released a document sum Internet Security Vulnerabilities. Thousands of org expanded Top-20 lists that followed one, two, and their efforts so they could close the most dangero servic these

This s additi Appli the e the p listed patch patch

We h Hence imme of lea gover vendors and consulting firms;

Sear Resp Opinio Should

High The ope that cou

The Busi Paid sea 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 sites as organic results.

The In securi can he and m

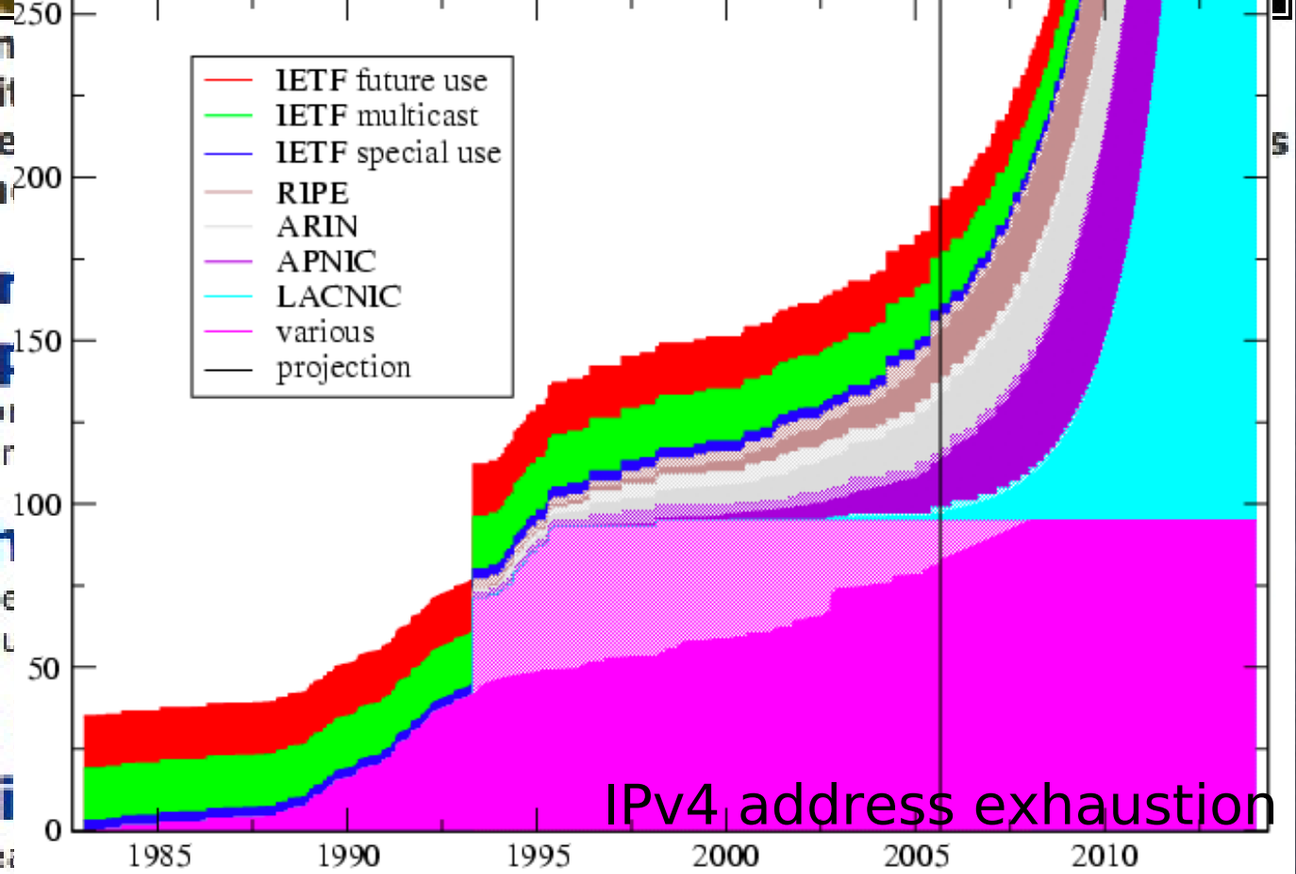

IPv4 address exhaustion

Home >> China

China adds top-level domain names

China's Ministry of Information domain name system in accor Regulations.

After the adjustment, ".MIL" wi "CN".

A new Internet domain name s

Under the new system, besi "NET" are temporarily set. It m servers under the managem Numbers (ICANN) of the Unite

Adobe Plugs Dreamweaver SQL Injection Flaw

How the internet killed the phone business

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

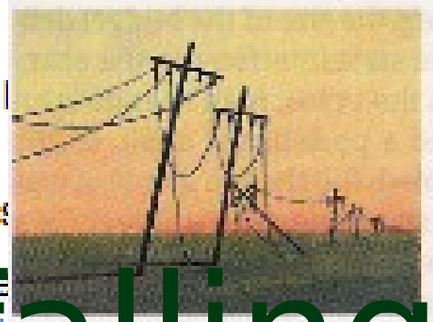

falling bits of sky

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. Thirty years ago, the personal computer was a classic example. It swept aside an 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 internet

International Telecommunication Union

Home : Office of the Secretary General : SPU

Next Generation Networks

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.)

Building Partnerships for Progress

All match

More

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

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

NETWORK WORLD

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 Blair Law, IDG News Service, 09/26/05

The latest buzzword in telecom isn't the name of a box, an application or a service. Instead, IMS is a way of organizing all those elements and more.

sundry "solutions"

in U.S. public sector resuming inquiry

DHS: data to validate security tools, SBGP, DNS
NIST: ways to measure DNSSEC penetration
DOE: way to estimate available bandwidth
FCC: way to measure outage
FTC: how to inform network neutrality debate
NCS/NSA: topology data for information assurance
GAO: cost of Internet katrina
NSF: can’t we just start over and do it right?

entire muni & community wireless networking movement….

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
er 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

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 have too little revenue, they cannot 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 transitions backbone to competitive market. no consideration of economics or security. kc proposes caida.org

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

soon, free markets 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 (seguridad, escalable, sostenible, adminstracion)

the 4 S’s

- *seguridad*: is the data toxic upon arrival?
- *escalable*: can we route/name/address earth’s needs?
- *sostenable*: is it economically viable?
- *administracion*: 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

network economics: dismal science(s)

known: economics of current architecture need study

has 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 offer help!

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.

cataloguing lessons

- although the Internet has over-achieved on plenty, it has underachieved on: security, scalability, sustainability, and stewardship. substantial oversights.
- our ability to measure is surprisingly abysmal, although policy history explains
- cooperative, data-sharing approaches to sound measurement and analysis are key to enlightened policy

dns traffic: queries to .cl server

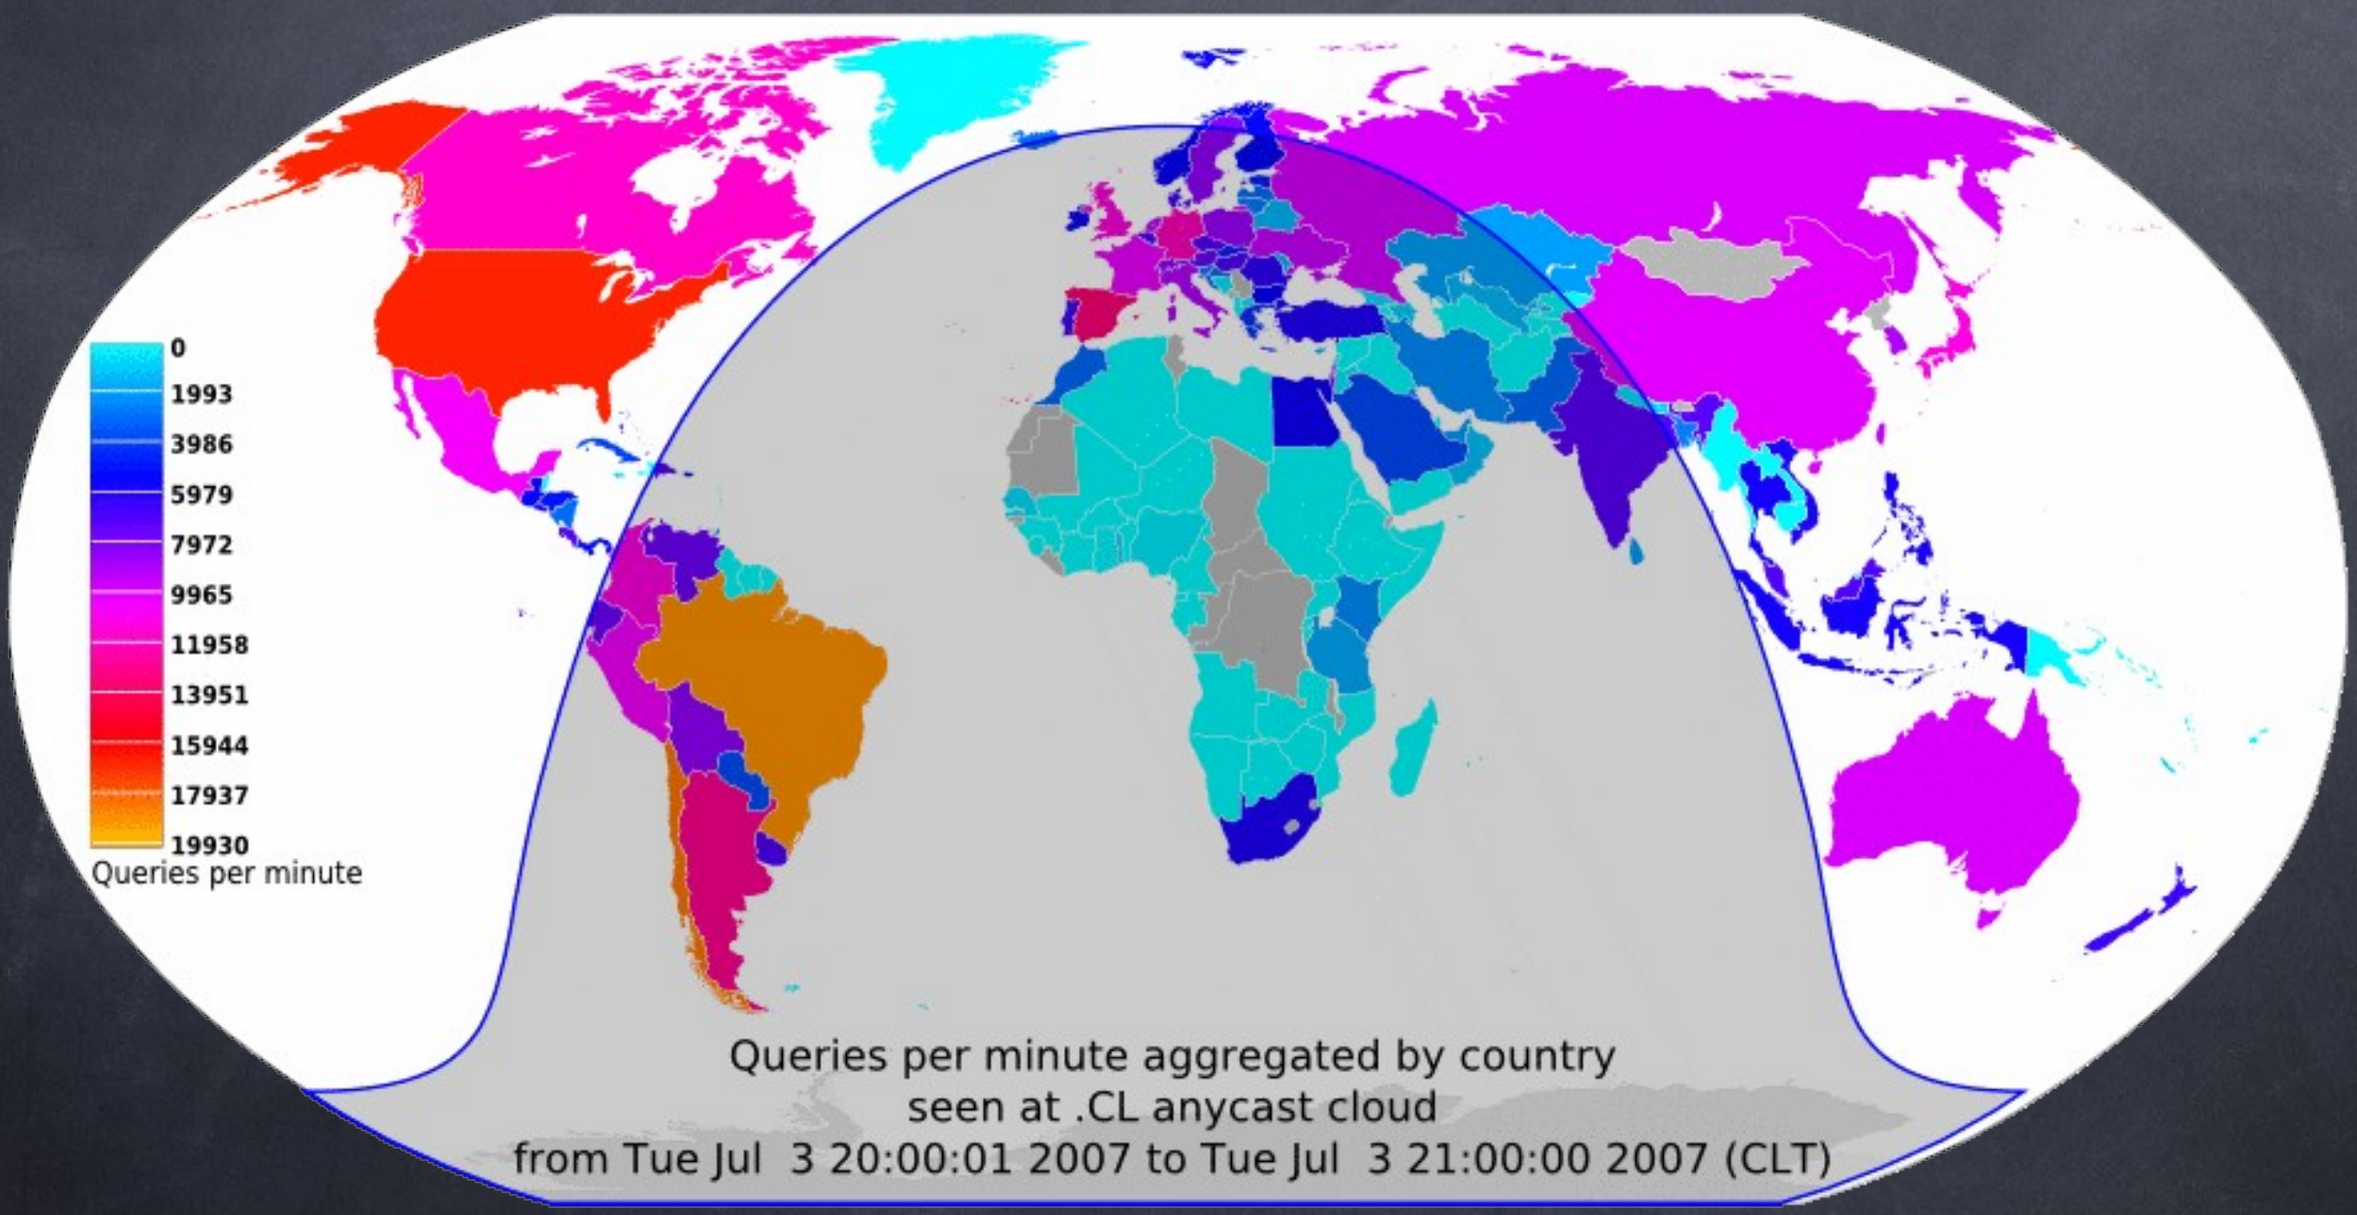

(Sebastian Castro of nic chile+caida)

active measurement: ***archipelago*** (ark)

- CAIDA's new measurement infrastructure
- 'operating system' for measurement
- launch 12 Sept 2007
- 28 active probers
- 5 are IPv6-capable
  - collaborators can run vetted measurements on security-hardened platform through simple API
  - general public can perform restricted measurements
  - support for meta-data mgt, analysis, and infoviz

IPv4 INTERNET
TOPOLOGY MAP
AS-level INTERNET GRAPH

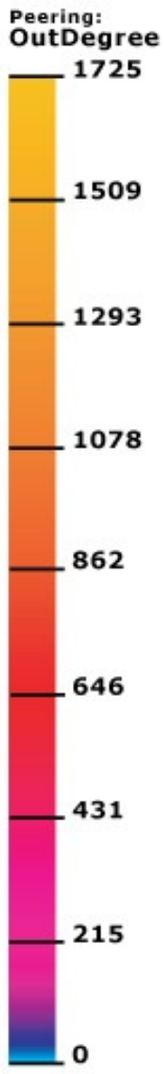

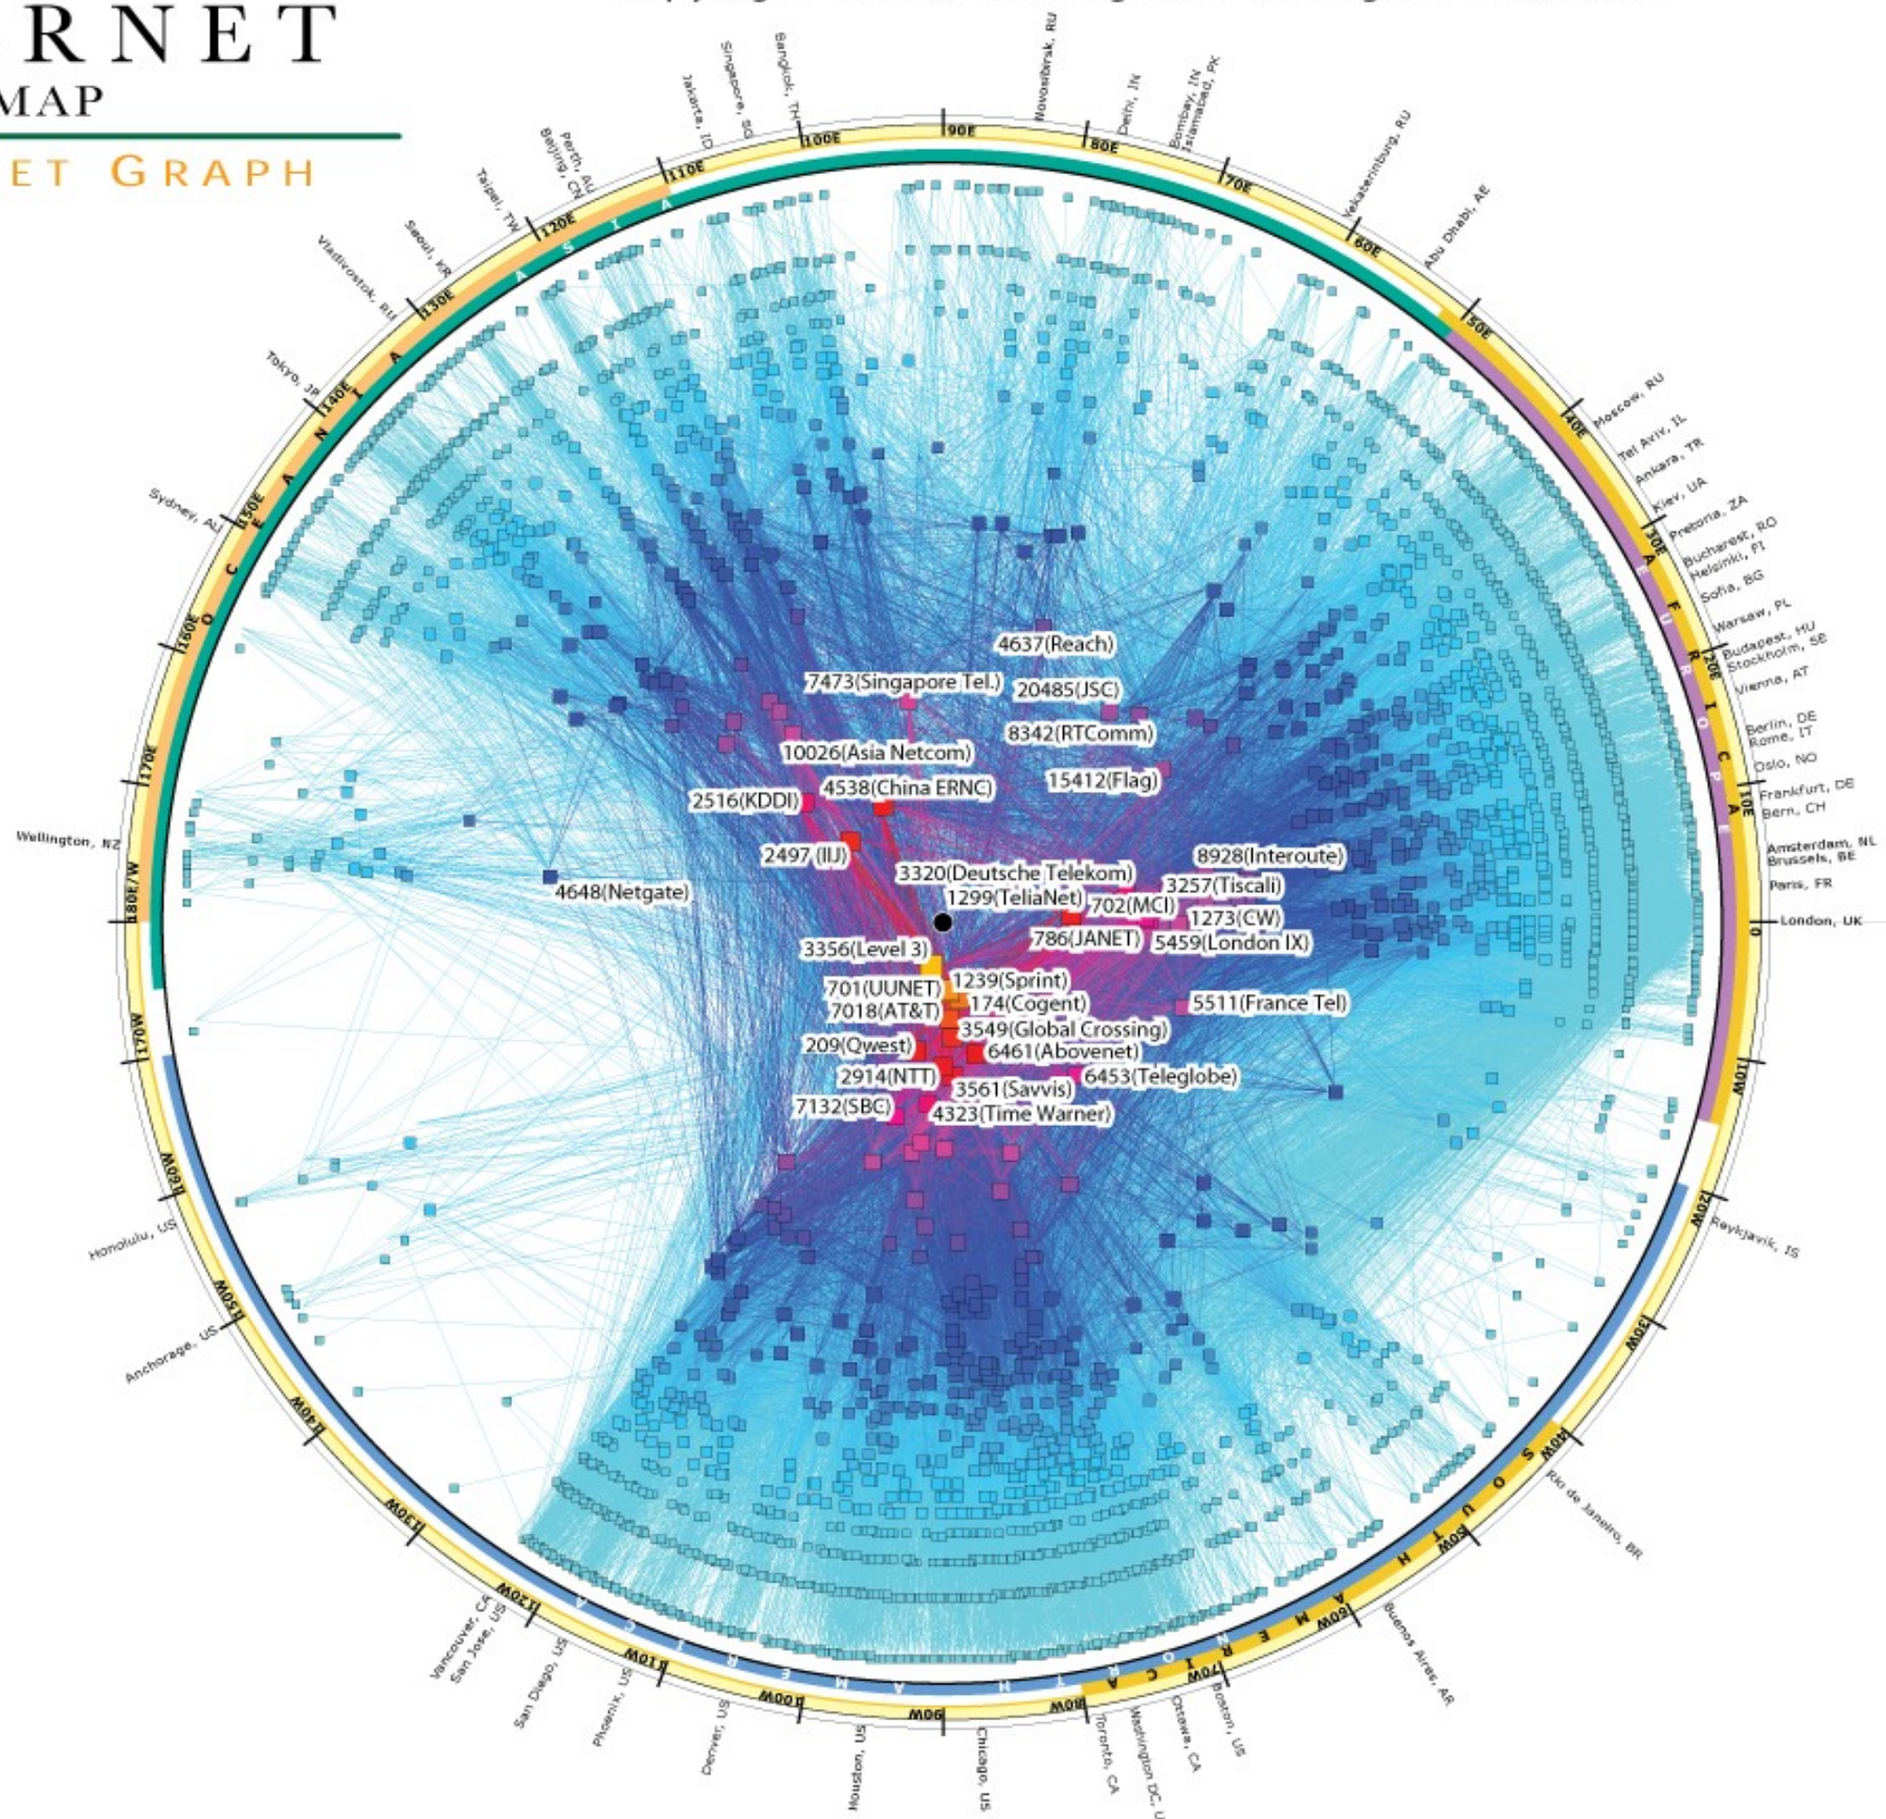

IPv6 INTERNET

TOPOLOGY MAP 2008

AS-level INTERNET GRAPH

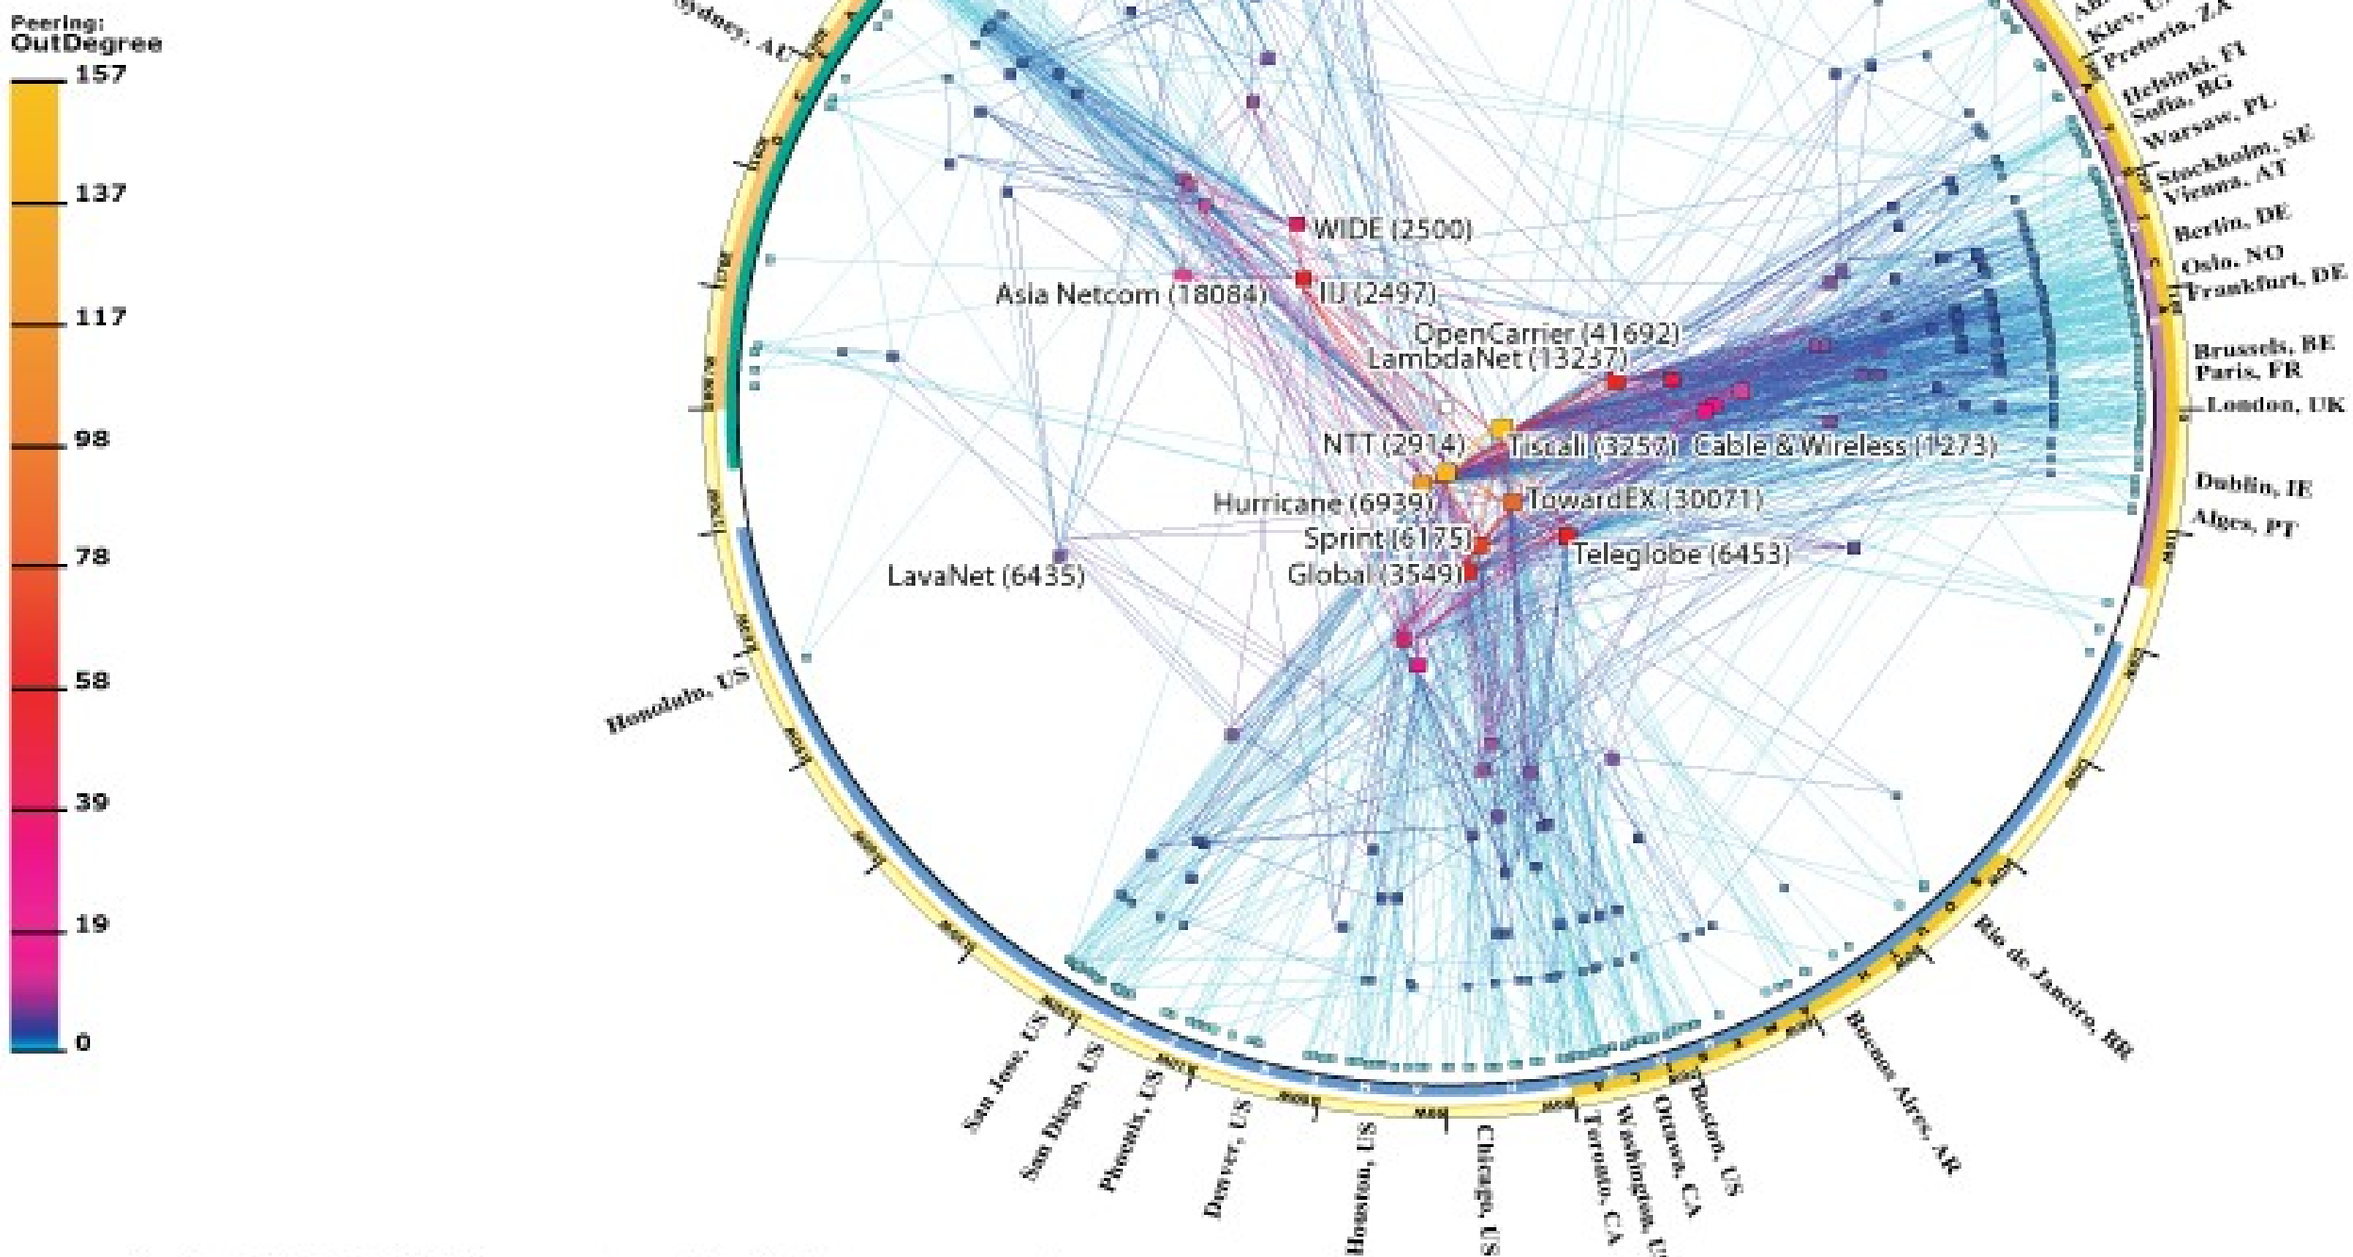

Internet policy: address exhaustion

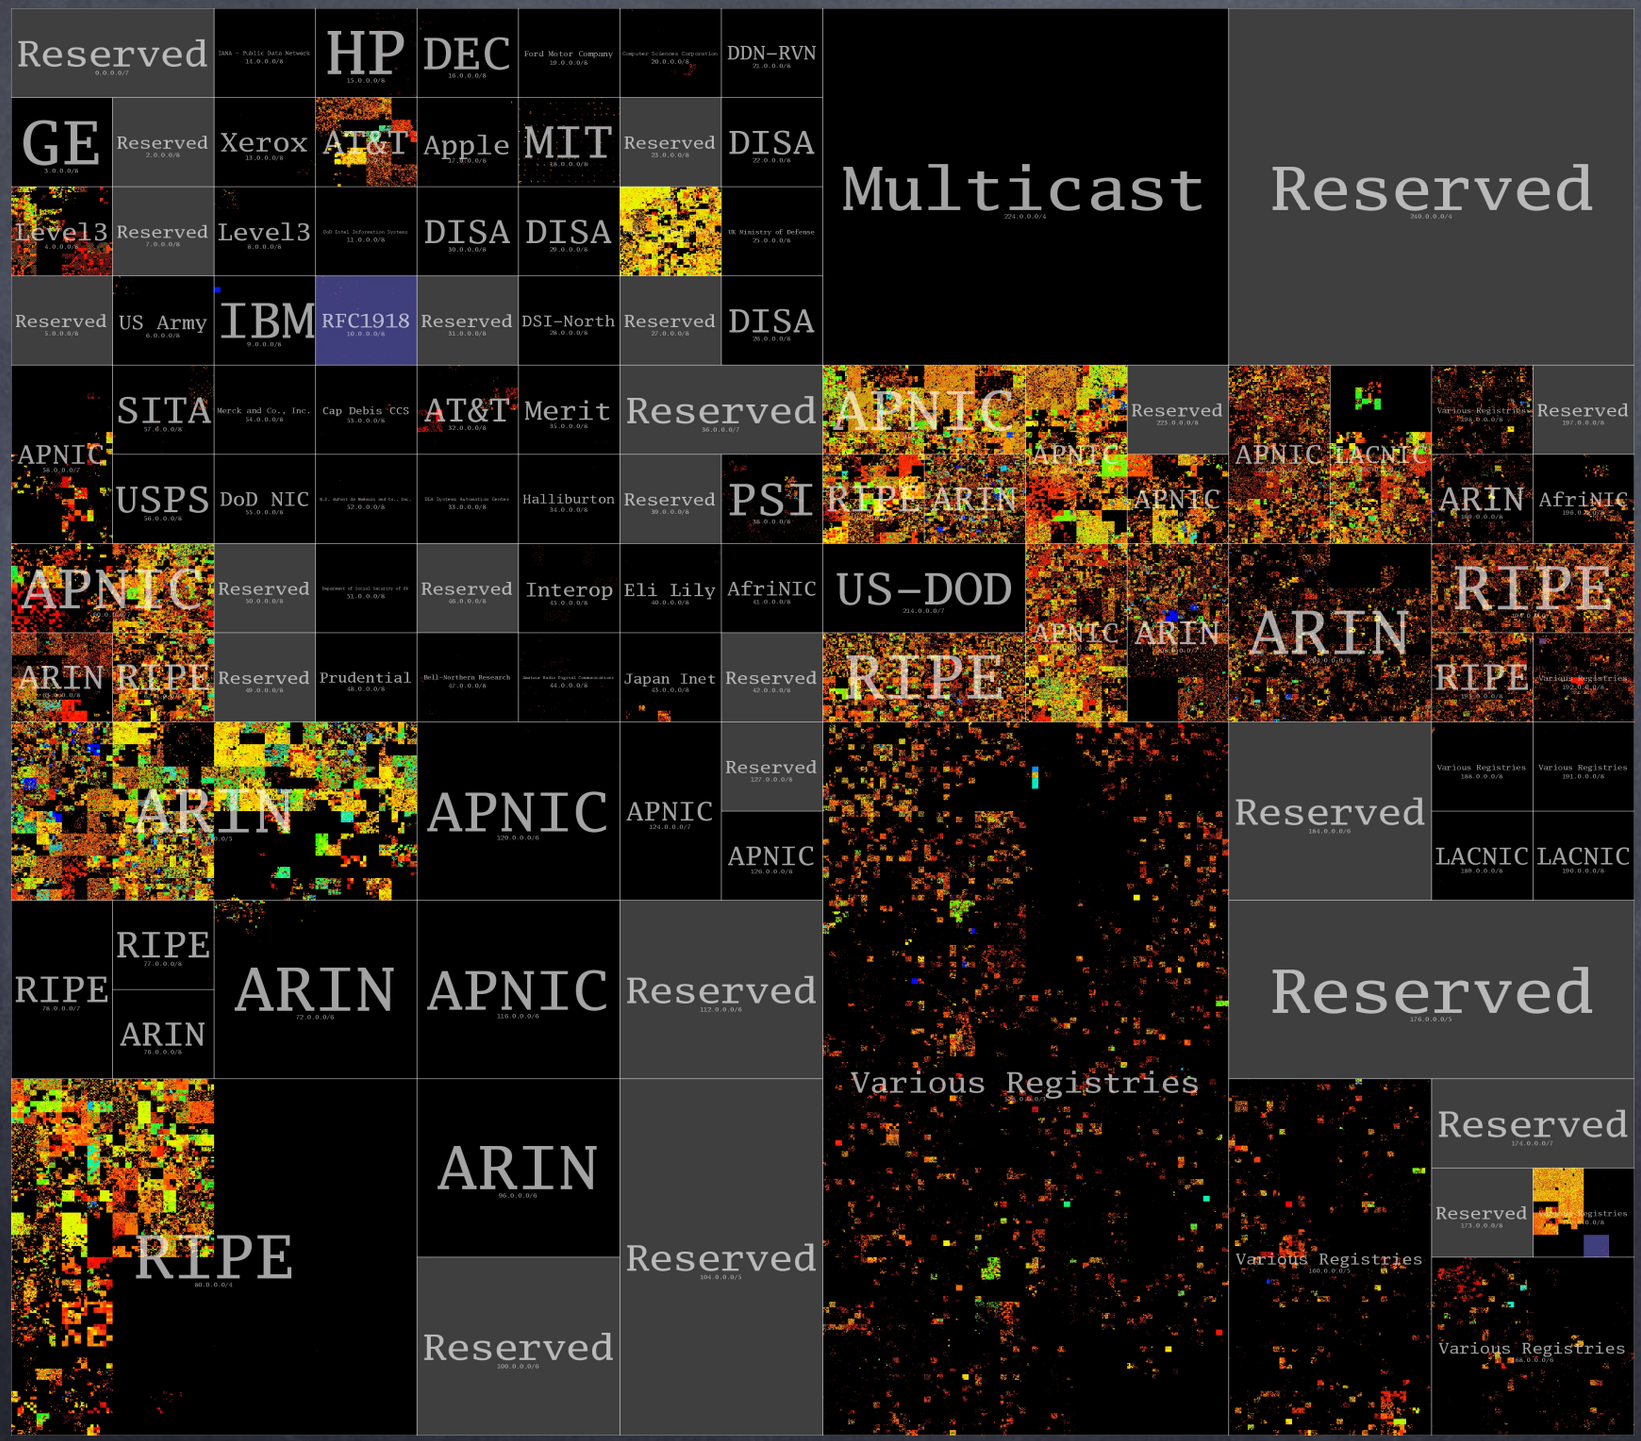

[ping data from isi.edu; poster by Duane Wessels@TMF]

Internet measurement data catalog

first catalog to support indexing and user annotations of Internet measurement data sets.

DatCat: (http://www.datcat.org)

facilitates searching for and sharing data among researchers,

enhances documentation of datasets via a public annotation system, and

advances network science by promoting reproducible research and persistent references.

CAIDA: summary of goals

Since 1997: narrowing the gap between Internet operations and science in face of global privatization

seek, analyze, communicate salient features of best available data on the Internet

forward-looking architectural research

navigate data-sharing challenges, by lowering technology barriers

support empirical needs of public sector

measurement accuracy is the only fail-safe means of distinguishing what is true from what one imagines, and even of defining what true means.

..this simple idea captures the essence of the physicist's mind and explains why they are always so obsessed with mathematics and numbers: through precision, one exposes falsehood.

a subtle but inevitable consequence of this attitude is that truth and measurement technology are inextricably linked.

-- robert b laughlin, a different universe,