Internet, Research, Stuff

Lars Eggert
lars@netapp.com

Outline

- Beginnings
- Style of Internet R&D
- The modern Internet
- Beyond

THE BEGINNINGS

I feel uncomfortable giving this talk in front of Vint, because he was there and I was in kindergarden

Lars Eggert
lars@netapp.com

Fostering Research was the Goal

Take any problem (...) and you find only a few people who can contribute effectively to its solution.

Those people must be brought into close intellectual partnership (...)

But bring these people together physically (...) and you have trouble, for the most creative people are often not the best team players (...)

Let them go their separate ways (...) and devote more time to the role of emperor than to the role of problem solver.

The principals still get together at meetings. (...) But the time scale of their communication stretches out (...) so that it may take a year to do a week's communicating.

There has to be some way of facilitating communication among people without bringing them together in one place.

The Computer as a Communication Device. J.C.R. Licklider, Robert W. Taylor. Science and Technology, April 1968.

And Research was the Foundation

1960s

- **Paul Baran**: message-block network

- **Don Davies**: packet-switching network

- **Leonard Kleinrock**: math theory

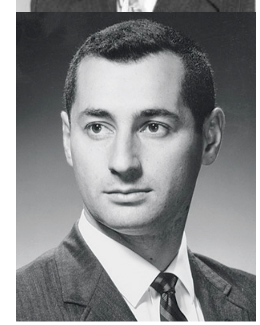

ARPANET

1962

- **J.C.R. Licklider** brought into DARPA to interconnect DoD computers
- Three terminals for SDC, UCB and MIT, wanted "one that goes anywhere"

1968

- Host level protocols development at UCLA under **Leonard Kleinrock** lead by **Steve Crocker**, with **Vint Cerf** and **Jon Postel**

1969 (Oct 29)

- First "LO" message between UCLA and SRI

ARPANET 1969-1976

http://internethalloffame.org/sites/default/files/08_28_2013_arpanet.gif
http://mercury.lcs.mit.edu/~jnc/tech/arpageo.html

Parallel & Complementary Efforts

- Michigan Educational Research Information Triad (**Merit**) Network (Michigan)
- **Mark I** at National Physics Laboratory (UK)
- **Louis Pouzin**: **CYCLADES** (France)

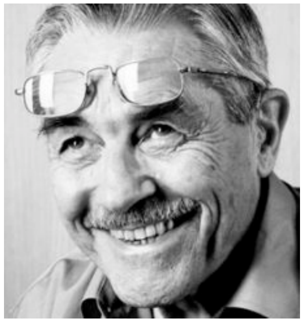

Kahn's Internetworking Ground Rules

- Each distinct **network** would have to **stand on its own** and **no internal changes** could be required (...) to connect it to the Internet
- Communications would be on a **best effort** basis (...) retransmitted from the source
- There would be **no information retained by the gateways** about the individual flows
- There would be **no global control** at the operations level

http://www.internetsociety.org/internet/what-internet/history-internet/brief-history-internet

Consequences, i.e., Early Research

- Algorithms to prevent lost packets from permanently disabling communications and enabling them to be successfully retransmitted from the source
- Providing for host-to-host “pipelining” so that multiple packets could be enroute from source to destination at the discretion of the participating hosts, if the intermediate networks allowed it
- Gateway functions to allow it to forward packets appropriately. This included interpreting IP headers for routing, handling interfaces, breaking packets into smaller pieces if necessary, etc.
- The need for end-end checksums, reassembly of packets from fragments and detection of duplicates, if any
- The need for global addressing
- Techniques for host-to-host flow control
- Interfacing with the various operating systems

http://www.internetsociety.org/internet/what-internet/history-internet/brief-history-internet

Consequences, i.e., Early Research

- Algorithms to prevent lost packets from permanently disabling
  Transport Protocol Design to be successfully
  retransmitted from the source
- Providing for host-to-host "pipelining" so that multiple packets
  Congestion Control rce to destination at the discretion of the
  participating hosts, if the intermediate networks allowed it
- Gateway functions to allow it to forward packets appropriately. This
  Packet Forwarding & Routing ting, handling interfaces,
  breaking packets into smaller pieces if necessary, etc.
- The need for end-end checksums, reassembly of packets from
  Transport Protocol Design tes, if any
- **Packet Forwarding** sing
- **Flow Control** t-to-host flow control
- **Stack Architecture** us operating systems

http://www.internetsociety.org/internet/what-internet/history-internet/brief-history-internet

Early Work Highlights (Subjective)

1973

- **Danny Cohen**: packet video/voice
- **Vint Cerf** and **Bob Kahn**: TCP/IP

1980s

- **Radia Perlman**: routing & spanning tree
- **Paul Mockapetris**: DNS (1983)
- **Van Jacobson**: TCP congestion control (1988)

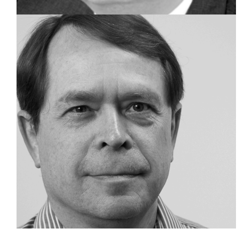

And much more: Email, FTP, NetNews, Telnet, …

The 1990s: Web & Telephony

1989

- **Tim Berners-Lee** creates WWW
- Publicly available in 1991

1993

- **Marc Andreessen** announces NCSA Mosaic browser

1996

- **Henning Schulzrinne** (co-)designs SIP, RTP, RTSP

The Internet Hourglass

Application layer
- Build apps with connections

Transport layer
- How to send over paths
- Connections

Network layer
- Network of networks
- Packets & network paths

Link layer
- Makes up a network

Physical layer
- Electrons, photons

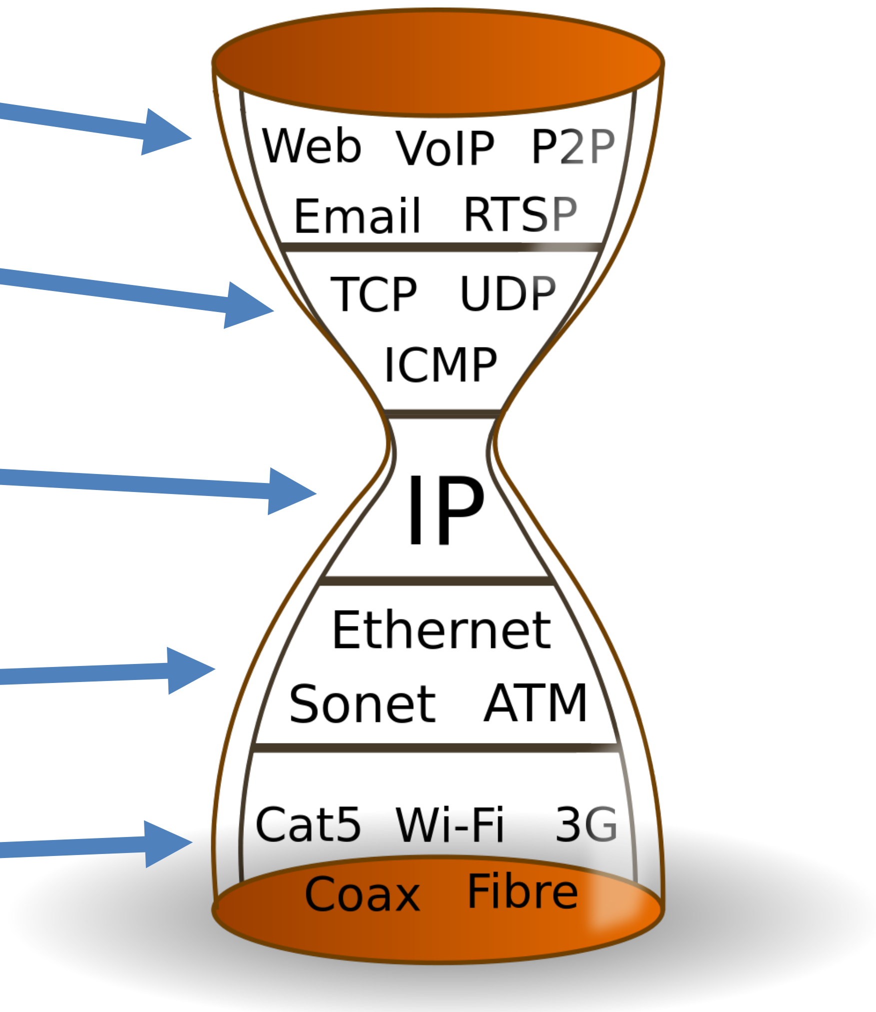

https://commons.wikimedia.org/wiki/File:Internet-hourglass.svg

The 1990s: Future-Proofing

- **IPv6** & **IPsec**
- **VPN** (secure tunnel) & later virtual nets
- **TCP**: models & performance work
- **IP** & link layer interactions (esp. wireless)
- Infrastructure work (**BGP** & **DNS**)
- Network address translation (**NAT**) & **firewalls**
- Building the web: **CDNs**, **SSL/TLS**, **HTTP 1.1**, …

The End of the Beginning

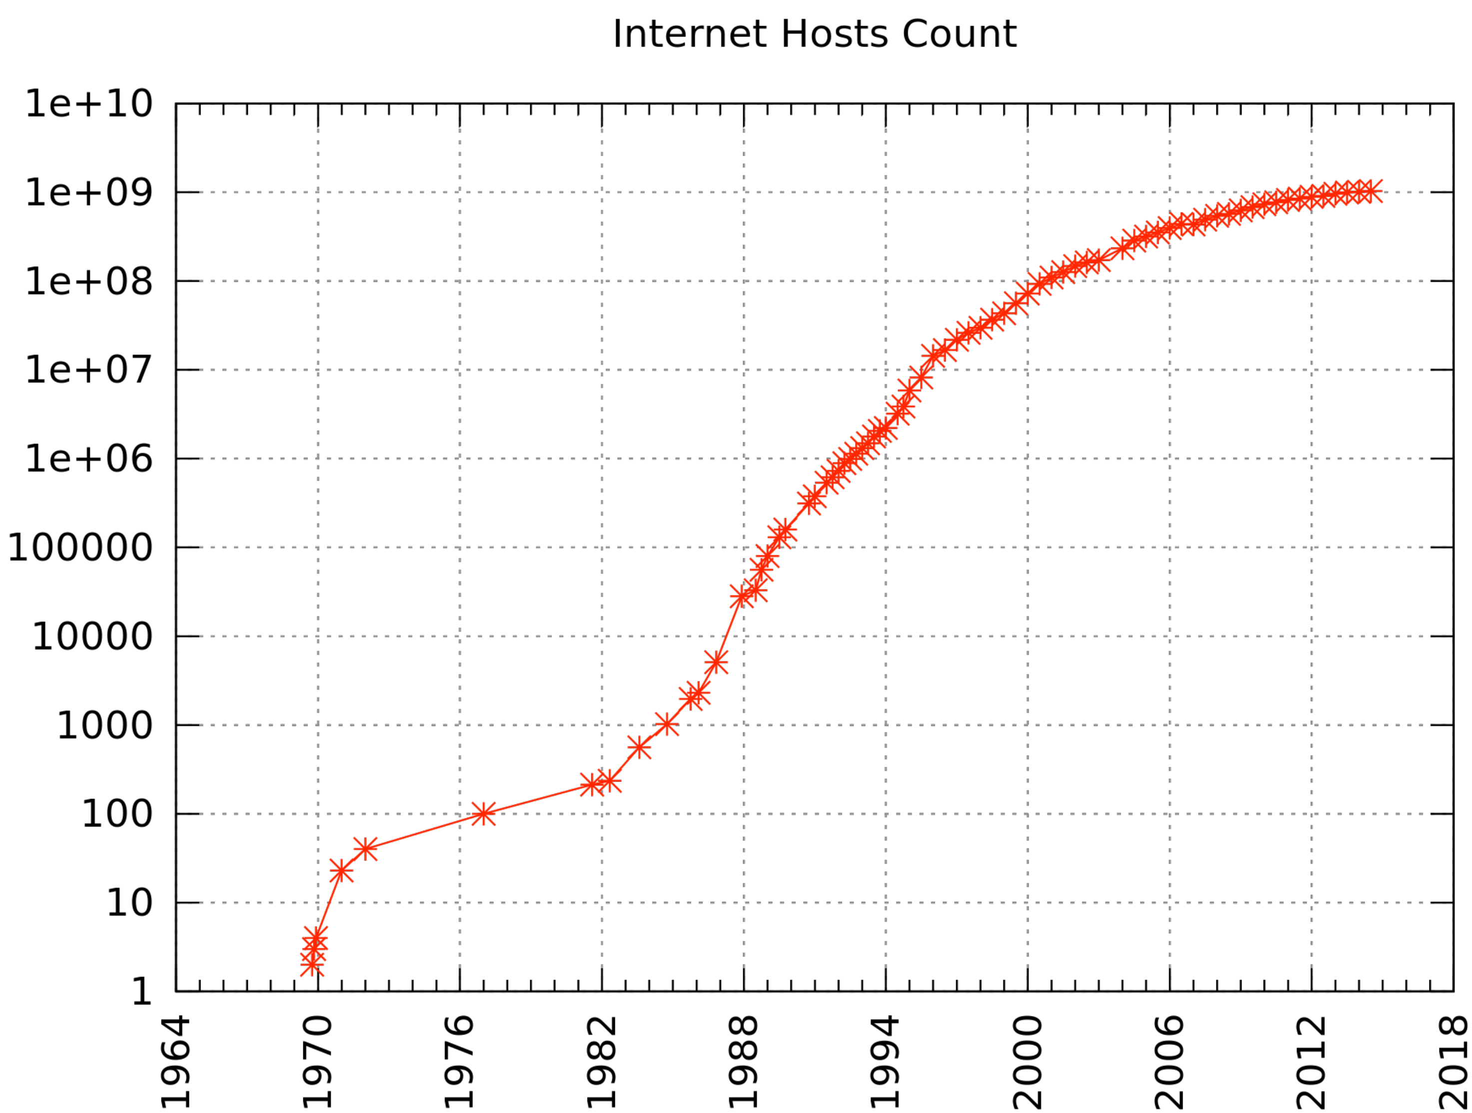

https://commons.wikimedia.org/wiki/File:Internet_Hosts_Count_log.svg

THE *STYLE* OF INTERNET R&D

“Open Access” from the Beginning

- **Request For Comments (RFC)** series created by **Steve Crocker** at UCLA in 1969
  - **Jon Postel** acting as editor (until 1998)
  - First distributed by (postal) mail, then FTP
- **Email** enabled open, wide-area collaboration
- **Source code** also freely shared
  - By email or netnews
  - Paving the way for open source later

Community Groups

1983

- Internet Configuration Control Board (**ICCB**) replaced by Task Forces, overseen by Internet Activities Board (**IAB**)

Late 1980s

- Explosive growth in Internet Engineering Task Force (**IETF**)
- Creation of Areas, Working Groups and Internet Engineering Steering Group (**IESG**); positioning of the IETF as standards body
- Creation of Internet Research Task Force (**IRTF**), holds other Task Forces

Early 1990s

- Formation of Internet Society (**ISOC**)
- Reorg of IAB into Internet Architecture Board under ISOC
- Reshuffling of relationship between IAB and IESG

Assigned Names & Numbers

- **Jon Postel** taking role of "czar of socket numbers" in 1972 (RFCs 349/433)
- Ended up coordinating all (?) Internet-related name and number spaces, eventually together with **Joyce Reynolds**
  - IP addresses, domain names, protocol parameters
- **Internet Assigned Numbers Authority (IANA)** first mentioned in RFC in 1990
- IANA function transferred to **ICANN** in Dec 1998

The Organizations Today

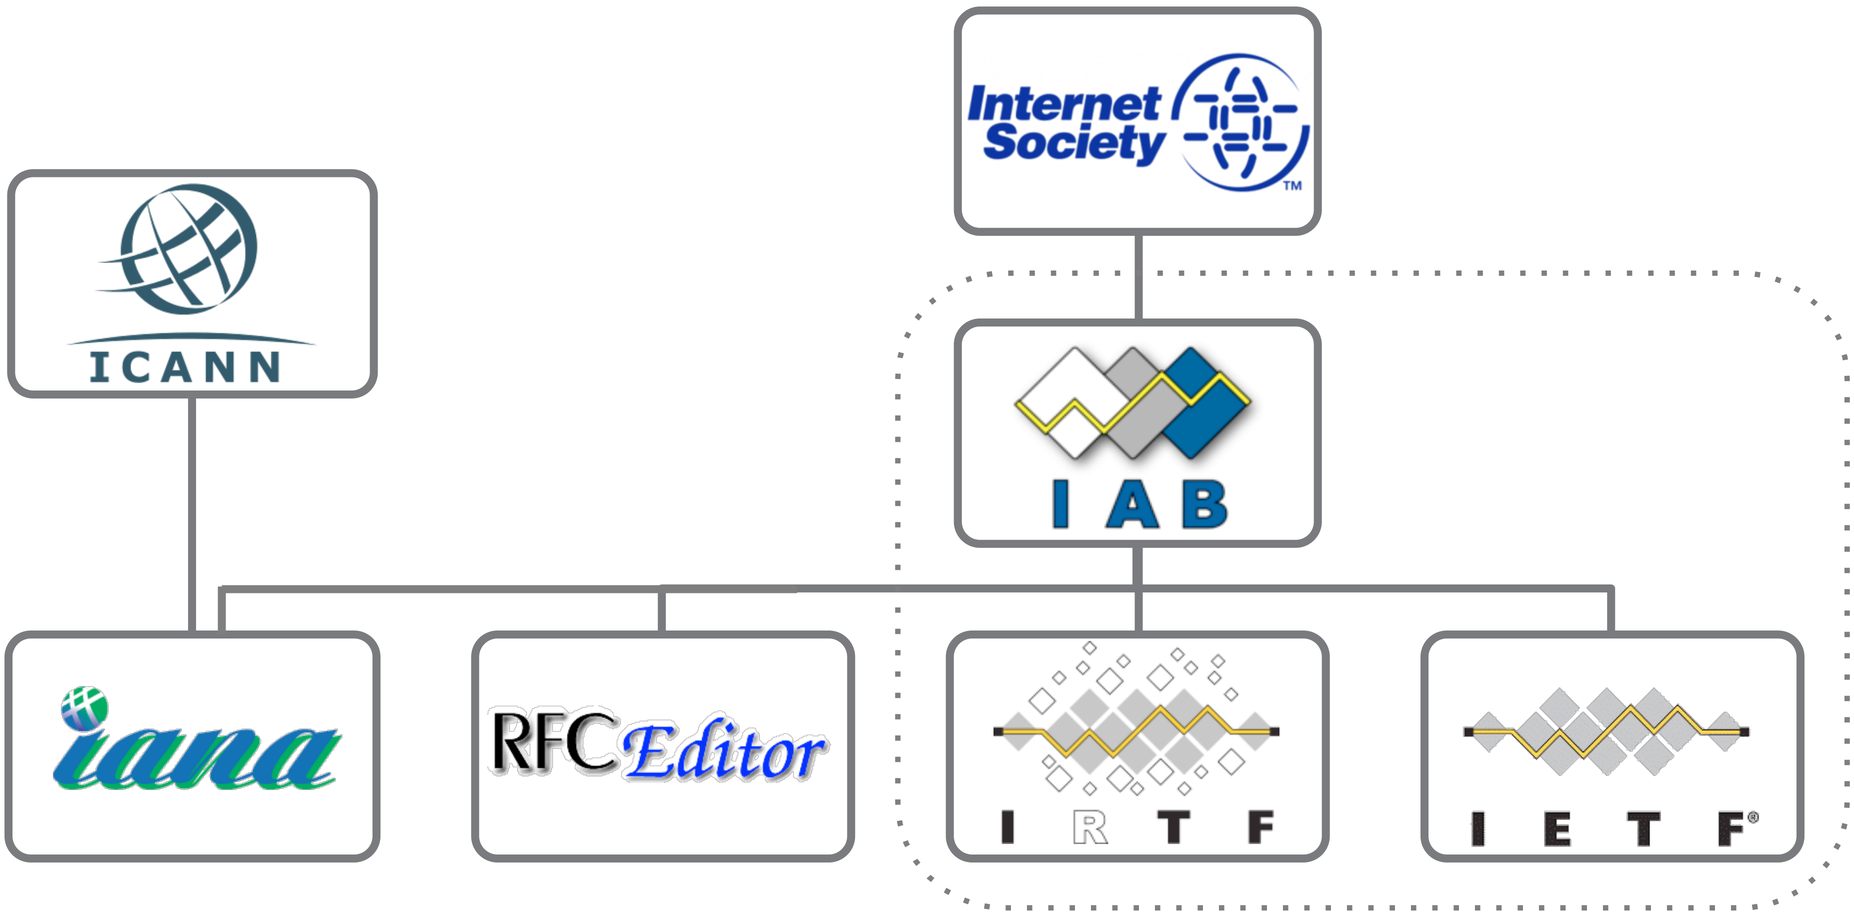

The Modern IETF

- The IETF is an **open**, **diverse** and **international** community
- Network **designers**, **operators**, **vendors**, **researchers**
- Common goal: **evolution of the Internet** architecture and protocols & **smooth operation** of the Internet
- **Participatory culture**; open to anyone
- **People**, not companies
- Produces Internet Standards & other docs
- It has a **research arm** – the IRTF

IETF by the Numbers

- **1-2K people** at 3 meetings/year
  - From **ca. 50 countries**
  - Many, many more on mailing lists
- Ca. **120 Working Groups**
- **8 Areas** with 15 Area Directors
- More than **7K RFCs** published
- **100K Internet-Draft revisions** submitted

- **IETF-90**, Toronto, Canada
  - 1175 people on site (152 newcomers)
  - 53 countries

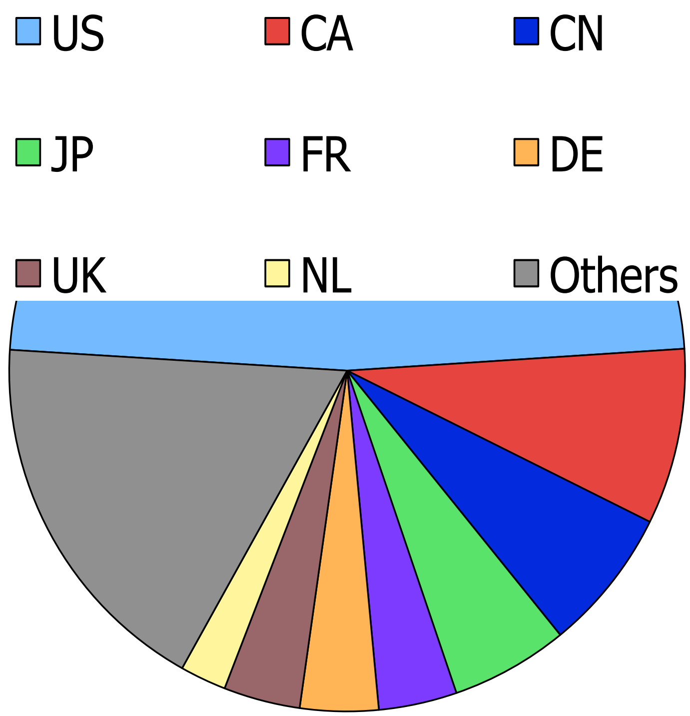

Most Active IETF Organizations

("at least one author with affiliation X")

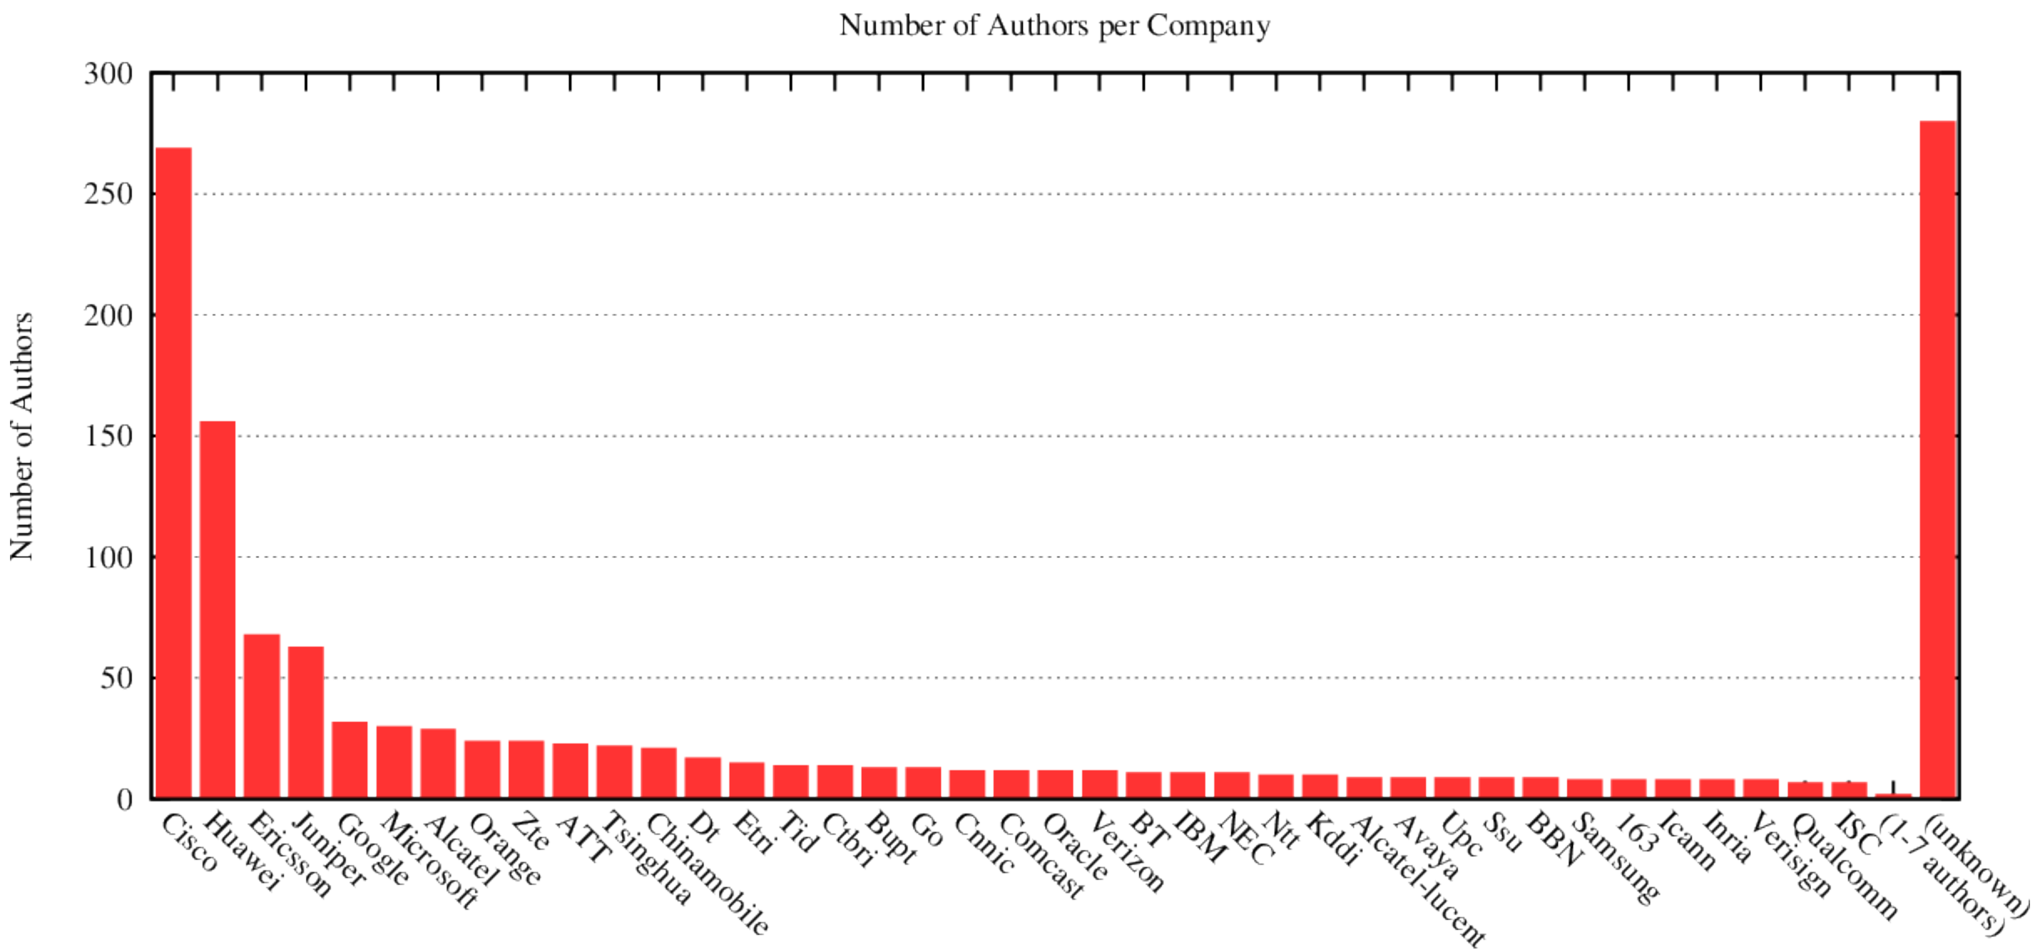

http://www.arkko.com/tools/stats/companydistr.html

Most Active Over Time

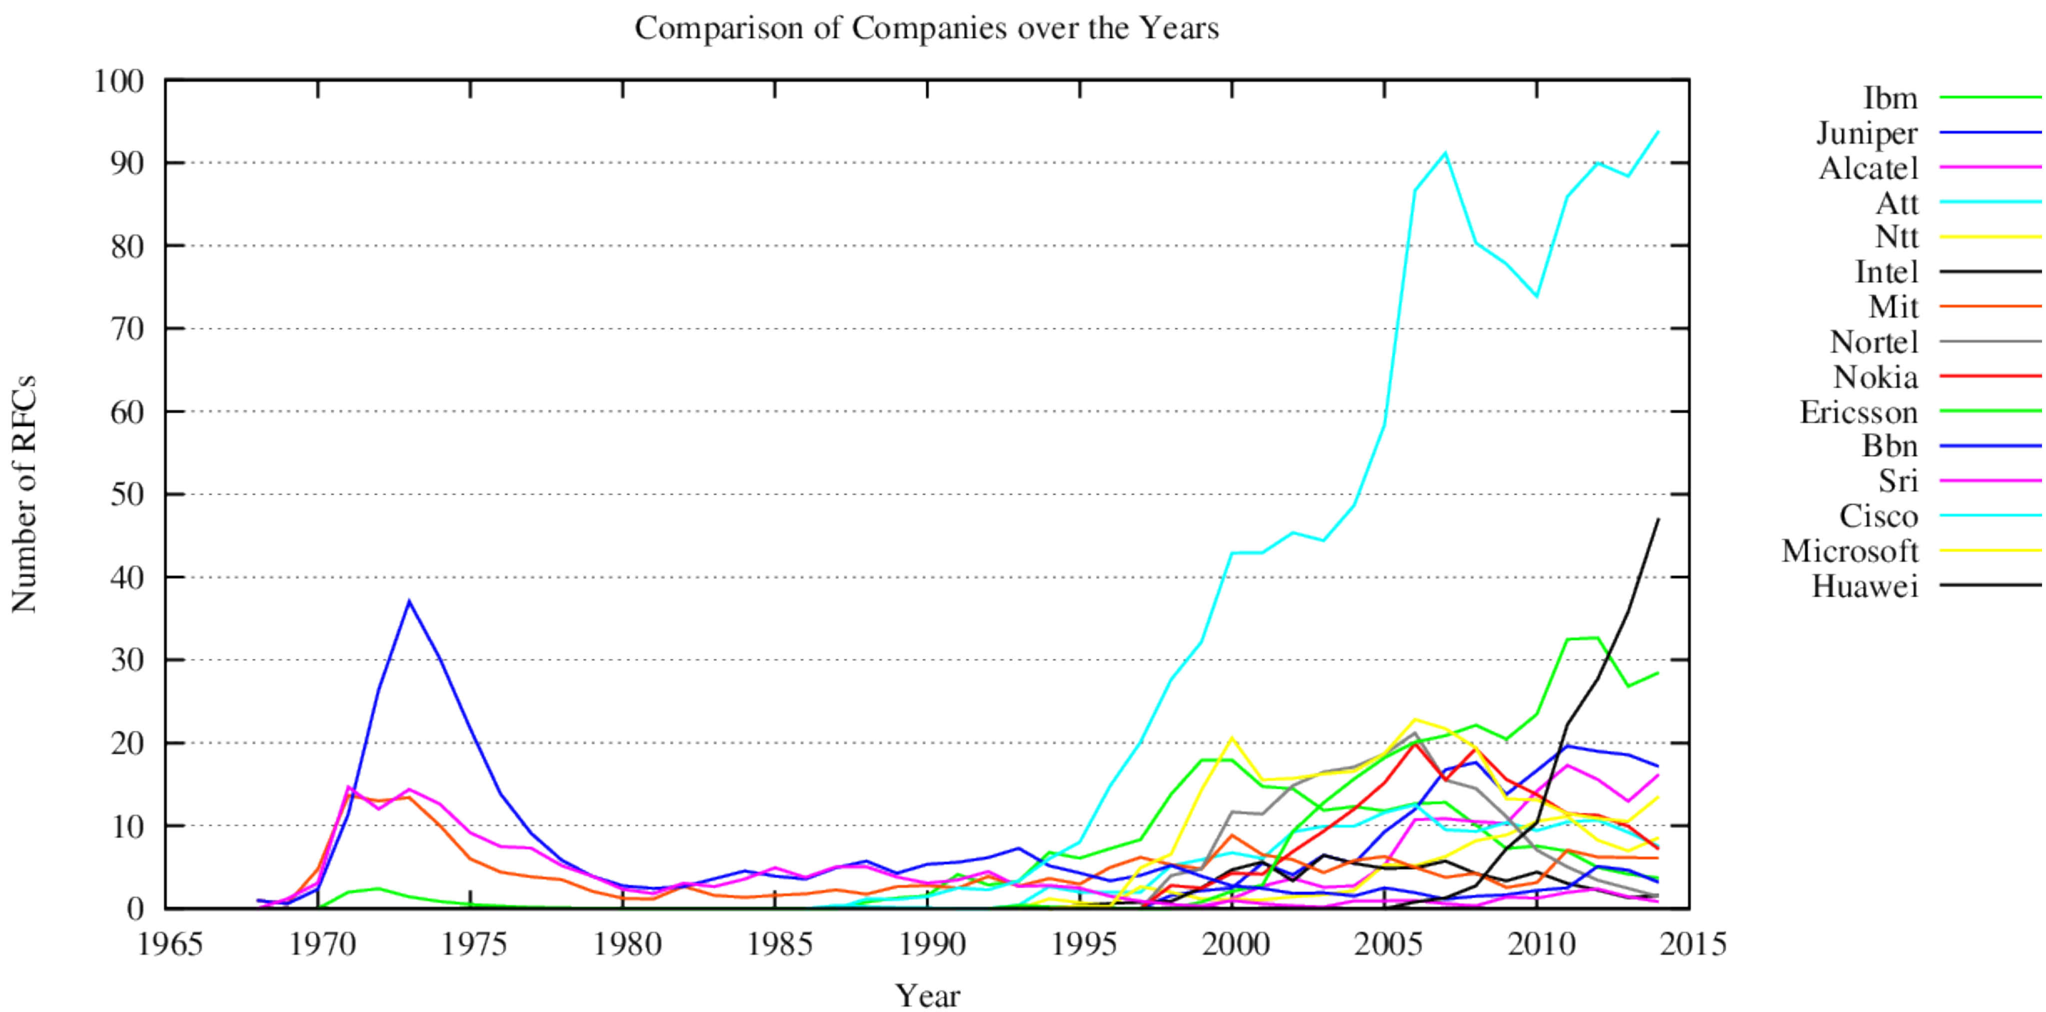

http://www.arkko.com/tools/rfcstats/companydistrhist.html

Authors with the Most Impact

26 for **Jon Postel** (210 RFCs cited 3904 times; avg. 18.59/RFC)

25 for **Keith McCloghrie** (95 RFCs cited 3116 times; avg. 32.80/RFC)

22 for **Steve Deering** (40 RFCs cited 1633 times; avg. 40.83/RFC)

20 for **Jonathan Rosenberg** (71 RFCs cited 1487 times; avg. 20.94/RFC)

19 for **Henning Schulzrinne** (87 RFCs cited 1838 times; avg. 21.13/RFC)

18 for **Marshall Rose** (68 RFCs cited 1421 times; avg. 20.90/RFC)

17 for **Yakov Rekhter** (77 RFCs cited 1238 times; avg. 16.08/RFC)

17 for **Russ Housley** (75 RFCs cited 976 times; avg. 13.01/RFC)

16 for **Joyce K. Reynolds** (64 RFCs cited 1086 times; avg. 16.97/RFC)

---

9 for **Vint Cerf** (43 RFCs cited 298 times; avg. 6.93/RFC)

http://www.arkko.com/tools/allstats/hindextop.html

Why Research Participation Matters

Researcher

- Learn about the **real problems**
- Work on meaningful issues – **help build the Internet**
- Understand what promotes and hinders **deployment**
- Meet potential **collaborators** and **funding** sources
- Have a realistic understanding of the **time commitments**

IETF

- Gains highly skilled, less-biased (?) **experts**
- Use academic results to create better **standards**
- Enable researchers to directly improve the Internet
- **Insight into trends** that impact standards down the road
- Accompany relevant topics in the IRTF research arm

IP
ON EVERYTHING

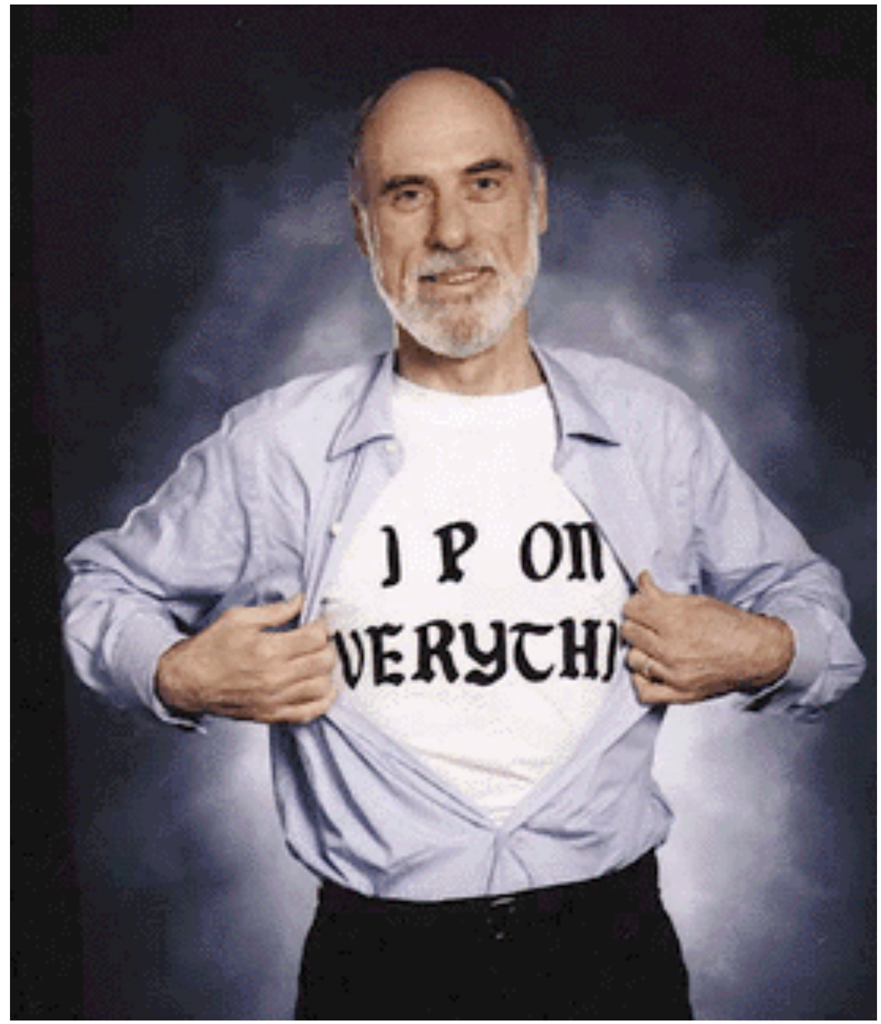

Mobile Networks

	1G	2G/2.5G	3G	4G	5G
Launch	1985	1992/1995	2002	2010	?
Bandwidth	2 Kbps	14.4-64 Kbps	2 Mbps	2-1000 Mbps	> 1Gbps
Technology	Analog cellular	Digital cellular	CDMA, **IP**	**IP**	**IP**
Service	Voice	2G: Voice, SMS 2.5G: + **data**	Audio, video, **data**	**Data**, audio, video	**Data** + etc.
Switching	Circuit	2G: Circuit 2.5G: + **packet**	**Packet** (circuit for air iface)	**Packet**	**Packet**
Core Network	PSTN	PSTN	Packet network	**Internet**	**Internet**

"IP is a Service" vs. "Services over IP"

- **Telcos had lots of misunderstanding about IP**
  - E.g., packet loss ratio as a metric
  - Desire to "optimize TCP" for wireless
- But biggest disagreement was about media
- Lots of bloat & waste making SIP, RTP, etc. interoperate with legacy telecom standards
  - Which the telcos are abandoning now, too
- **Left us with an ugly pile of technology that is begging to be pushed aside**

The Rise of the Datacenter

- Server farms to push content for growing web
  - Driven by commodity compute & network gear
- Enabled solving of previously hard problems
  - With old & new distributed systems techniques
  - At massive scale and interactive response times
- This amplified datacenter growth
  - **Google @ ONS 2015: 1Pb/s bisection bandwidth** (i.e., 100K servers @ 10Gb/s)
- All (?) running IP on technology (?)

Datacenter Challenges

- **Low, predictable latencies are critical**
- IP data plane generally OK
- IP control plane needed some engineering
  - Esp. because of the desire to virtualize
- TCP needed some serious research
  - Was optimized over decades for throughput
  - Forgotten interactions with queuing resurfaced
- **Continued pressure to “optimize TCP/IP”**
  - Should we?

Internet of Things

- Connecting a myriad of tiny, embedded sensors and actuators to the Internet
  - **Everything with a microcontroller will be online**
- IETF distinguishes device classes by KB RAM/KB flash (RFC 7288)
  - **Class 2:** 50/250 (<< Rasp Pi) **small full IP stack**
  - **Class 1:** 10/100 **constrained IP stack**
  - **Class 0:** << class 1 **IP difficult**

$5

Class 1

STM32W108HB STM32W108CB STM32W108CC STM32W108CZ

High-performance, IEEE 802.15.4 wireless system-on-chip with up to 256 Kbyte of embedded Flash memory

Features

- Complete system-on-chip
  - 32-bit ARM® Cortex®-M3 processor
  - 2.4 GHz IEEE 802.15.4 transceiver and lower MAC
  - 128/192/256-Kbyte Flash, 8/12/16-Kbyte RAM memory
  - AES128 encryption accelerator
  - Flexible ADC, SPI/UART/I^{2}C serial communications, and general-purpose timers
  - 24 highly configurable GPIOs with Schmitt trigger inputs
- Industry-leading ARM® Cortex®-M3 processor
  - Leading 32-bit processing performance
  - Highly efficient Thumb®-2 instruction set
  - Operation at 6, 12 or 24 MHz
  - Flexible nested vectored interrupt controller
- Low power consumption, advanced management
  - Receive current (w/ CPU): 27 mA
  - Transmit current (w/ CPU, +3 dBm TX): 31 mA
  - Low deep sleep current, with retained RAM and GPIO: 400 nA/800 nA with/without sleep timer

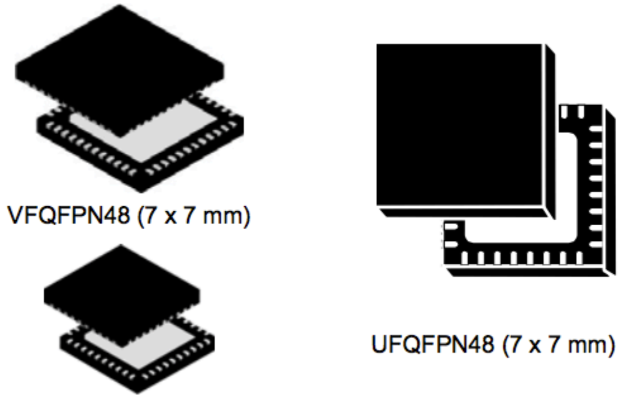

VFQFPN48 (7 x 7 mm)

UFQFPN48 (7 x 7 mm)

VFQFPN46 (6 x 6 mm)

  - Robust WiFi and Bluetooth coexistence
- Innovative network and processor debug
  - Non-intrusive hardware packet trace
  - Serial wire/JTAG interface
  - Standard ARM debug capabilities: Flash patch and breakpoint; data watchpoint and trace; instrumentation trace macrocell
- Application flexibility
  - Single voltage operation: 2.1-3.6 V with internal 1.8 V and 1.25 V regulators
  - Optional 32.768 kHz crystal for higher timer accuracy
  - Low external component count with single 24 MHz crystal
  - Support for external power amplifier
  - Small 7x7 mm 48-pin VFQFPN and

(Some) IoT Challenges

- Operate on **µW**= nodes mostly sleep
- Network bandwitdh ~100Kb/s
- Radio environment unstable
- Physical packet size limits (~100B)
- Infrequent ping(-pong) message exchanges

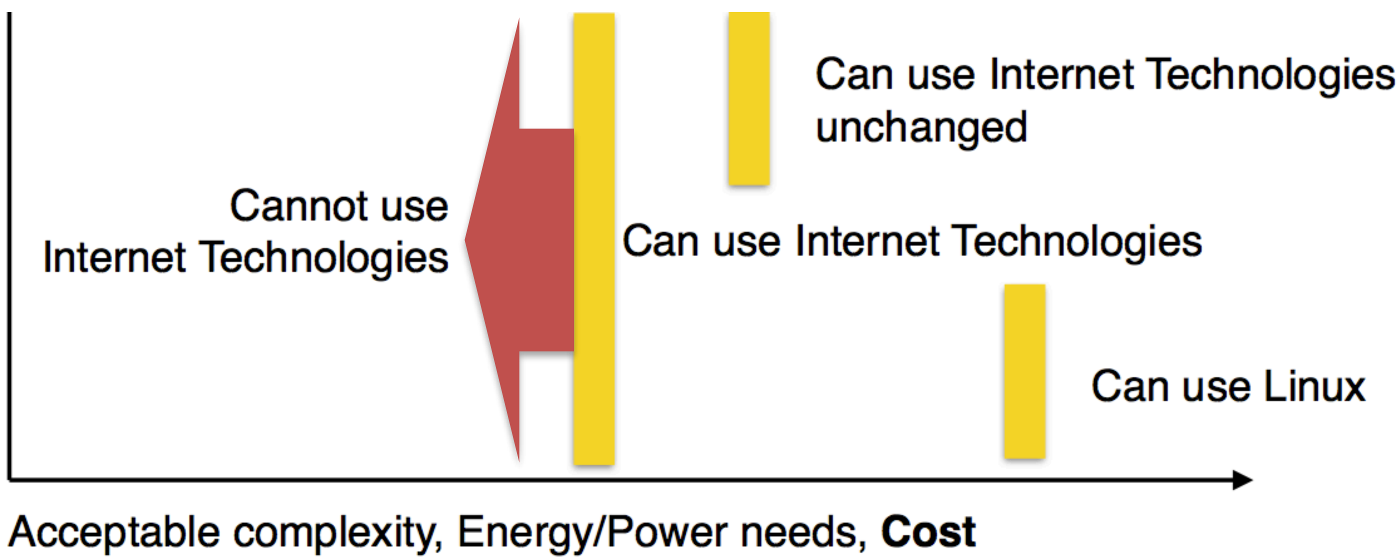

Figure adapted from Carsten Bormann

Interplanetary Internet

- No contemporaneous end-to-end path may ever exist between communicating nodes

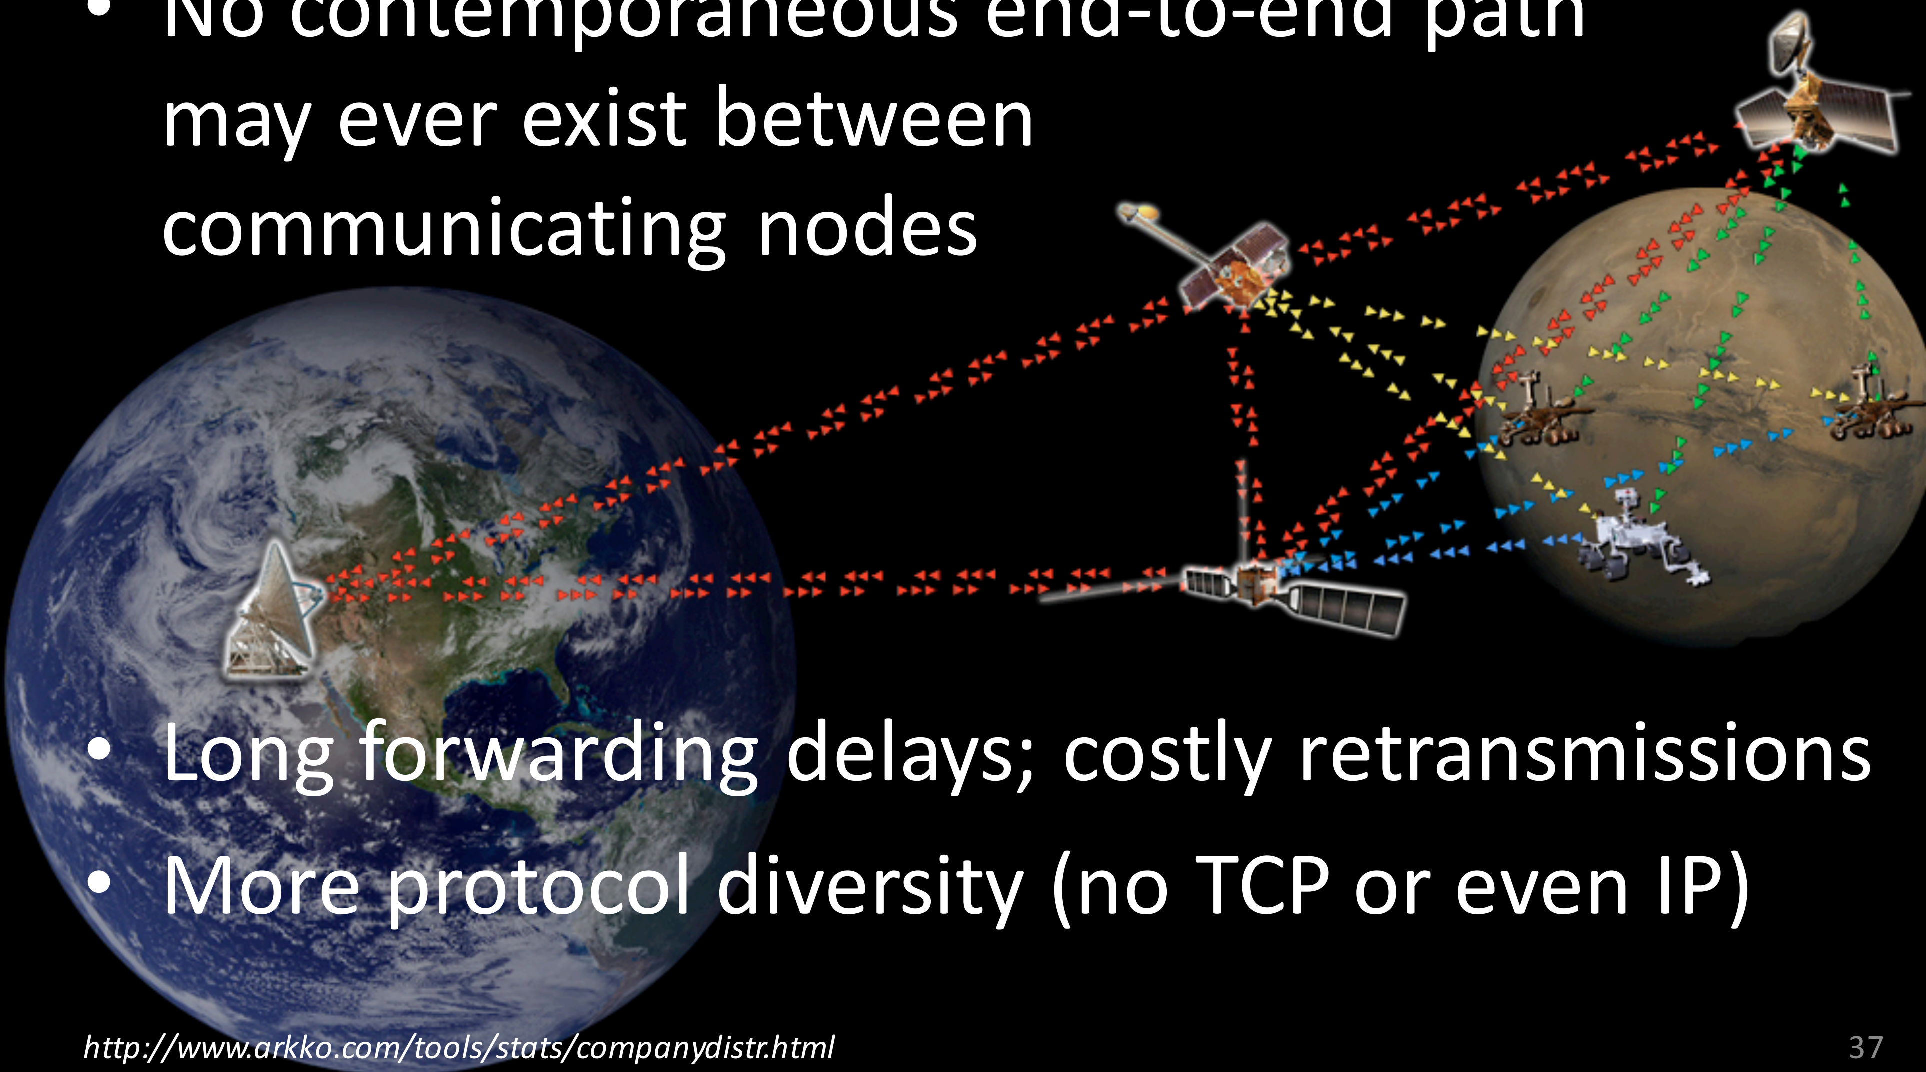

- Long forwarding delays; costly retransmissions
- More protocol diversity (no TCP or even IP)

http://www.arkko.com/tools/stats/companydistr.html

SCaN Notional Integrated Communication Architecture

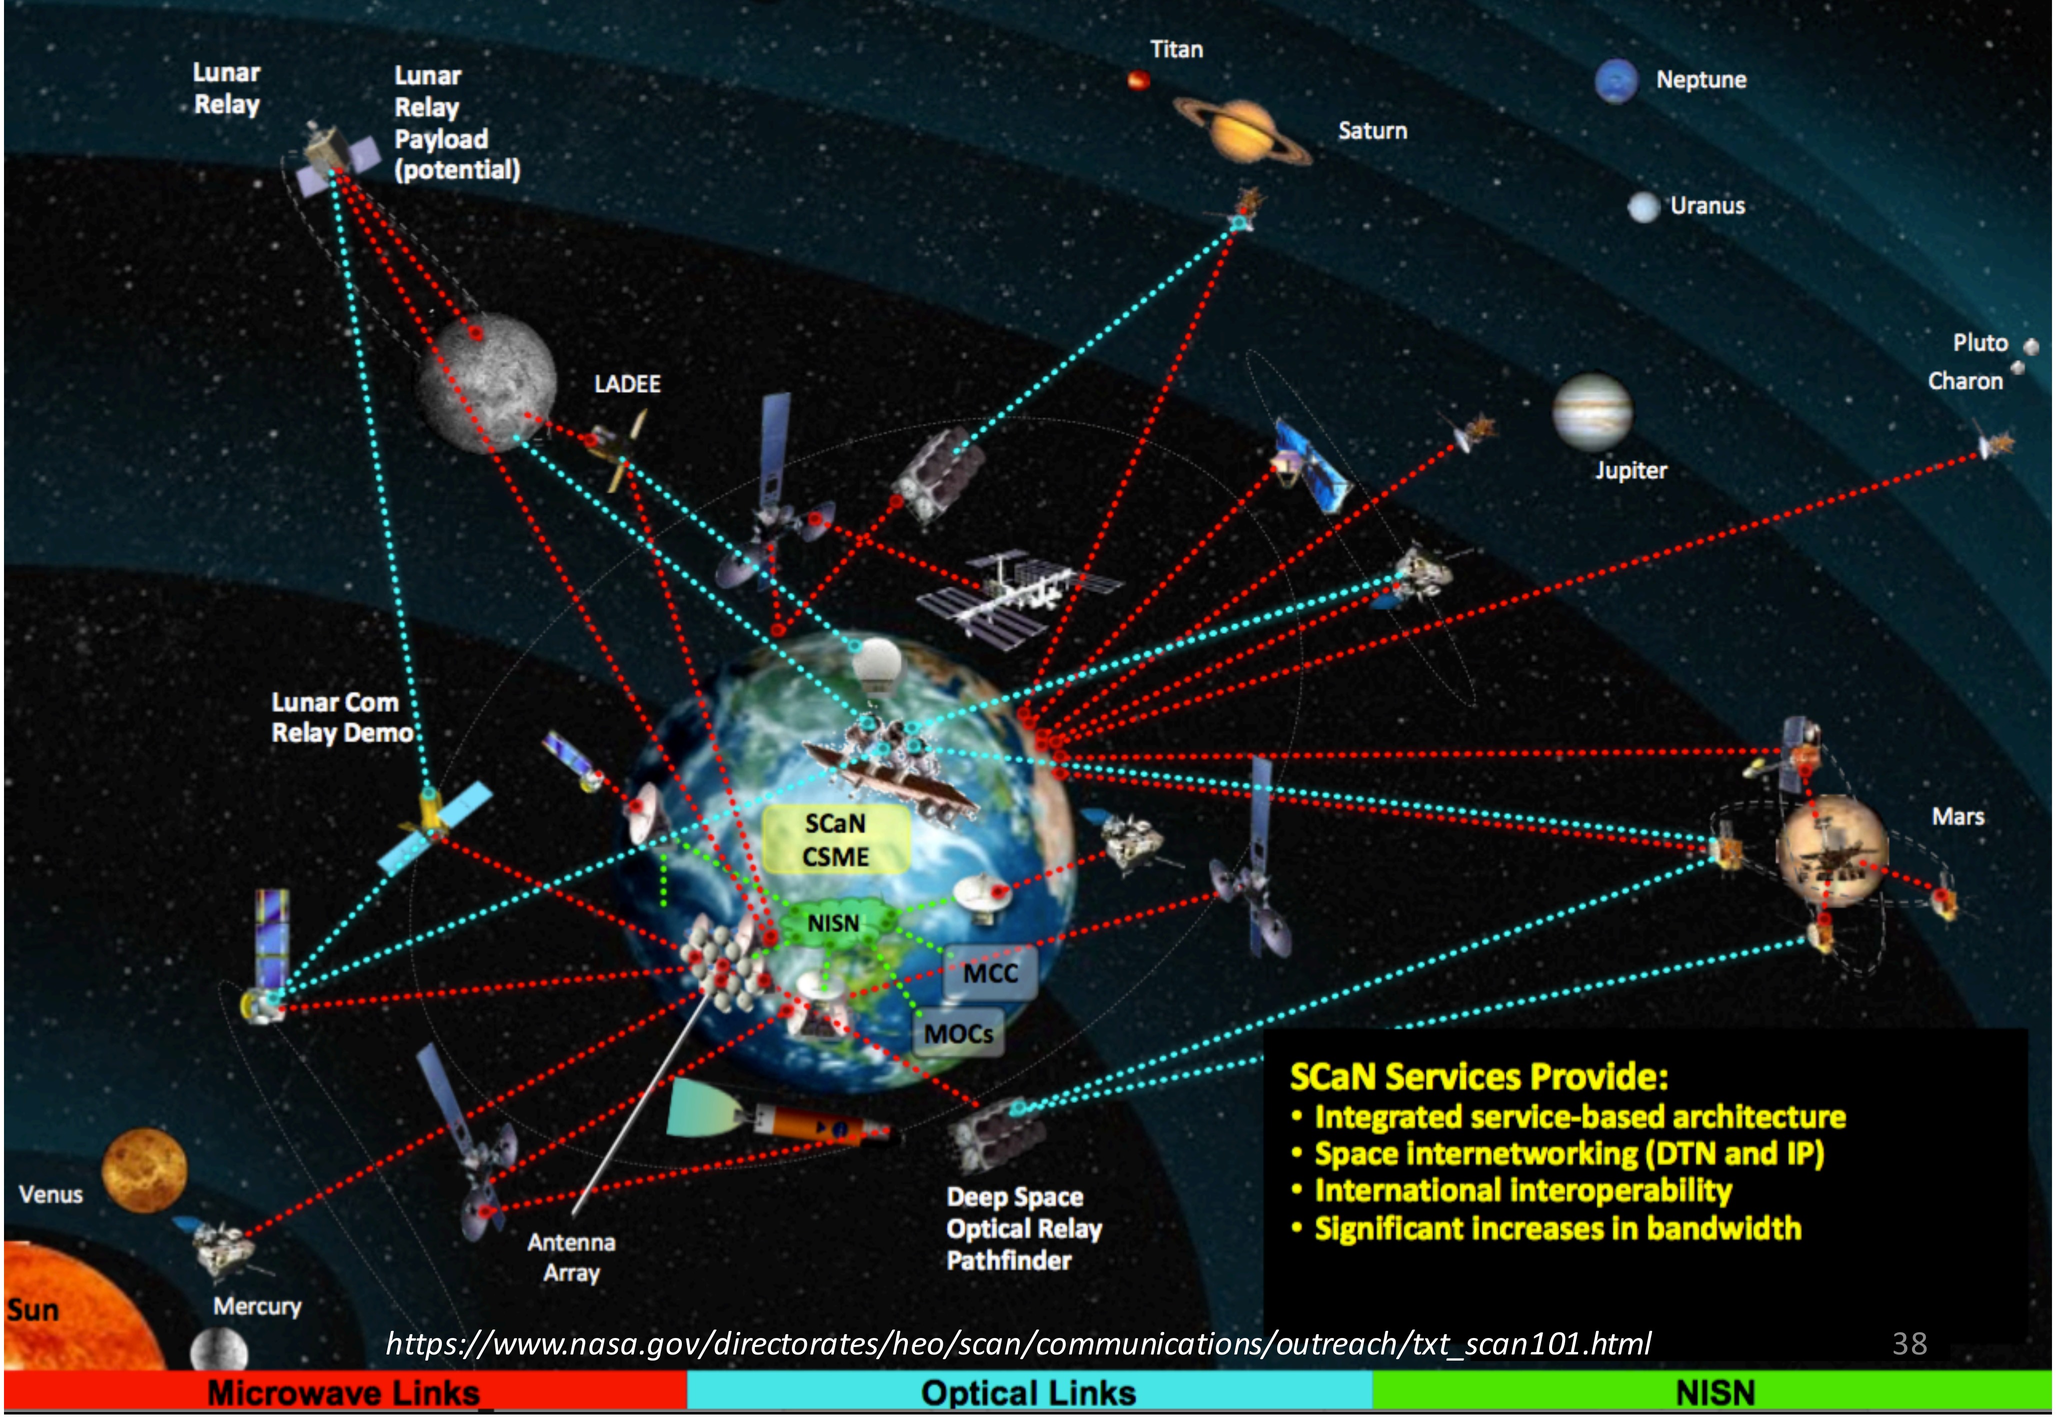

SCaN Services Provide:
- Integrated service-based architecture
- Space internetworking (DTN and IP)
- International interoperability
- Significant increases in bandwidth

https://www.nasa.gov/directorates/heo/scan/communications/outreach/txt_scan101.html

ASK VINT DURING THE BREAK!

THE FUTURE

Internet Everywhere?

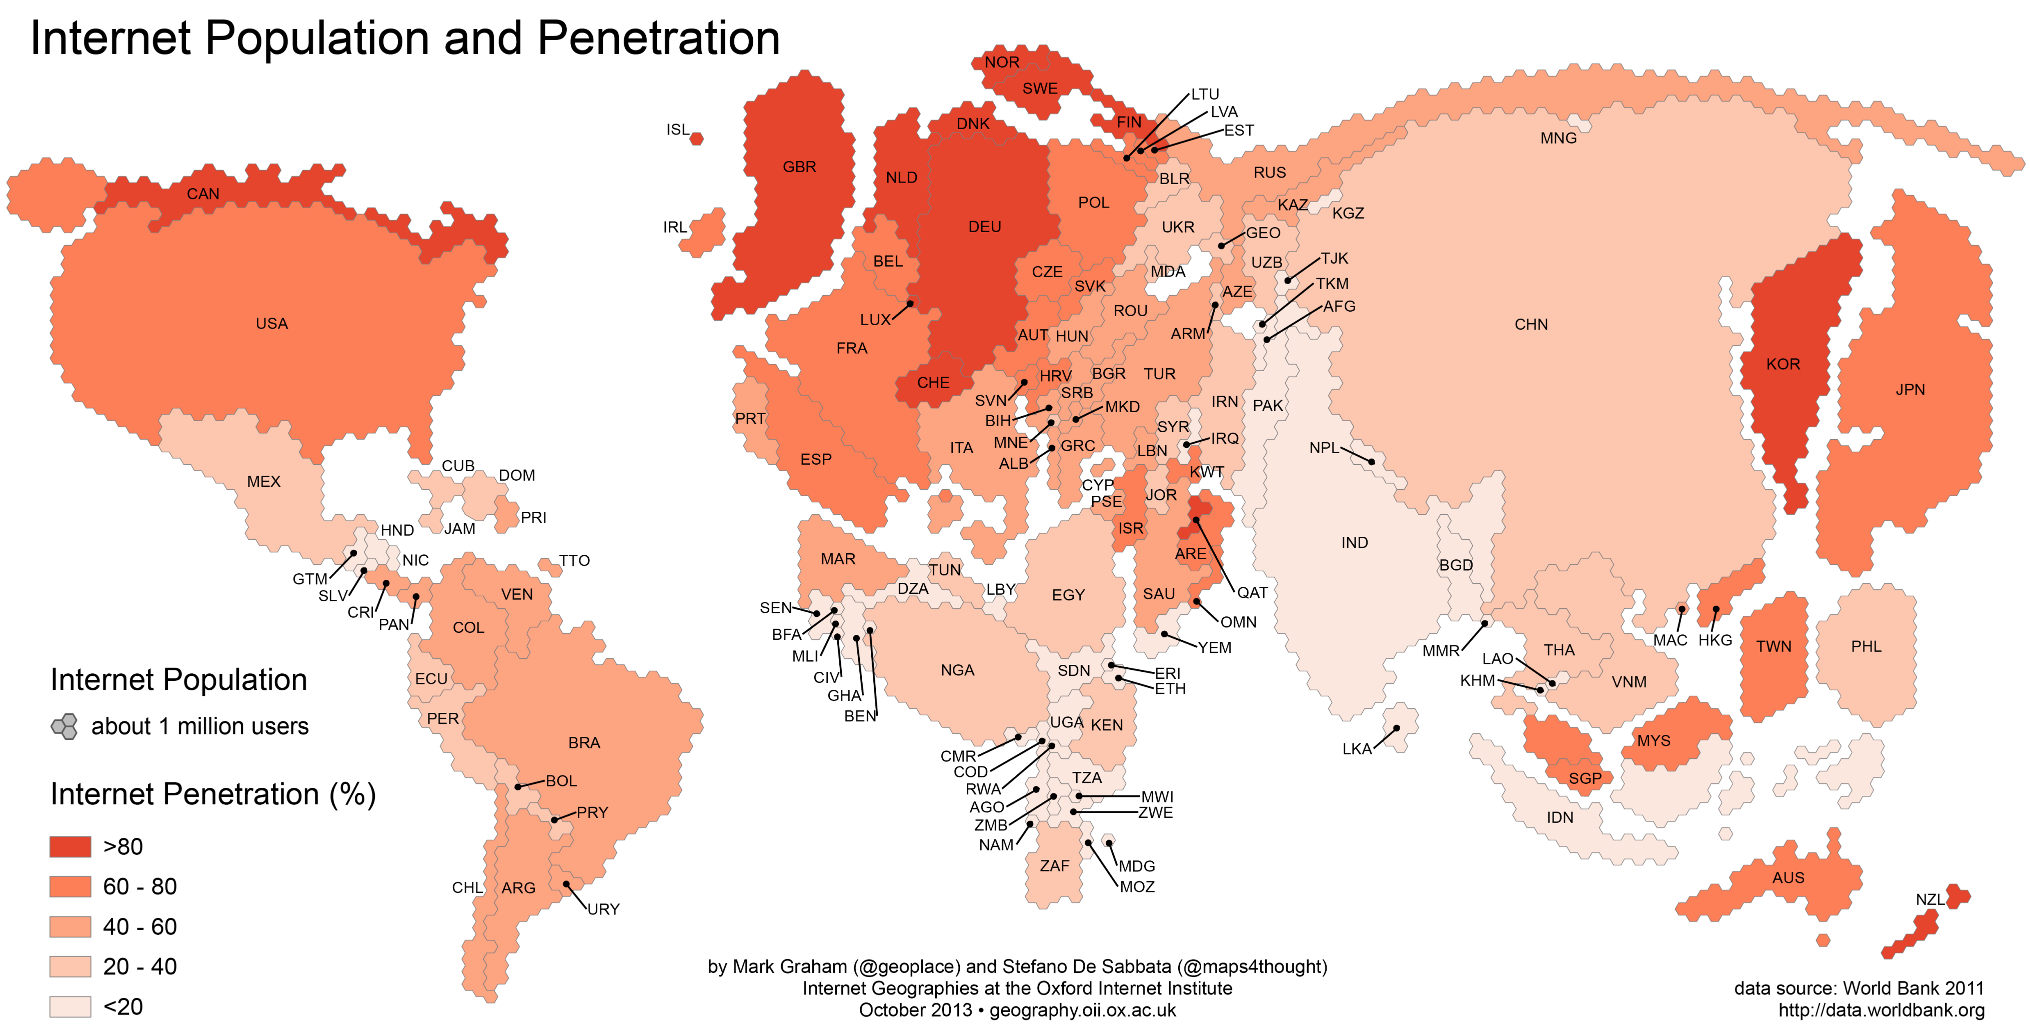

Broadband Affordability

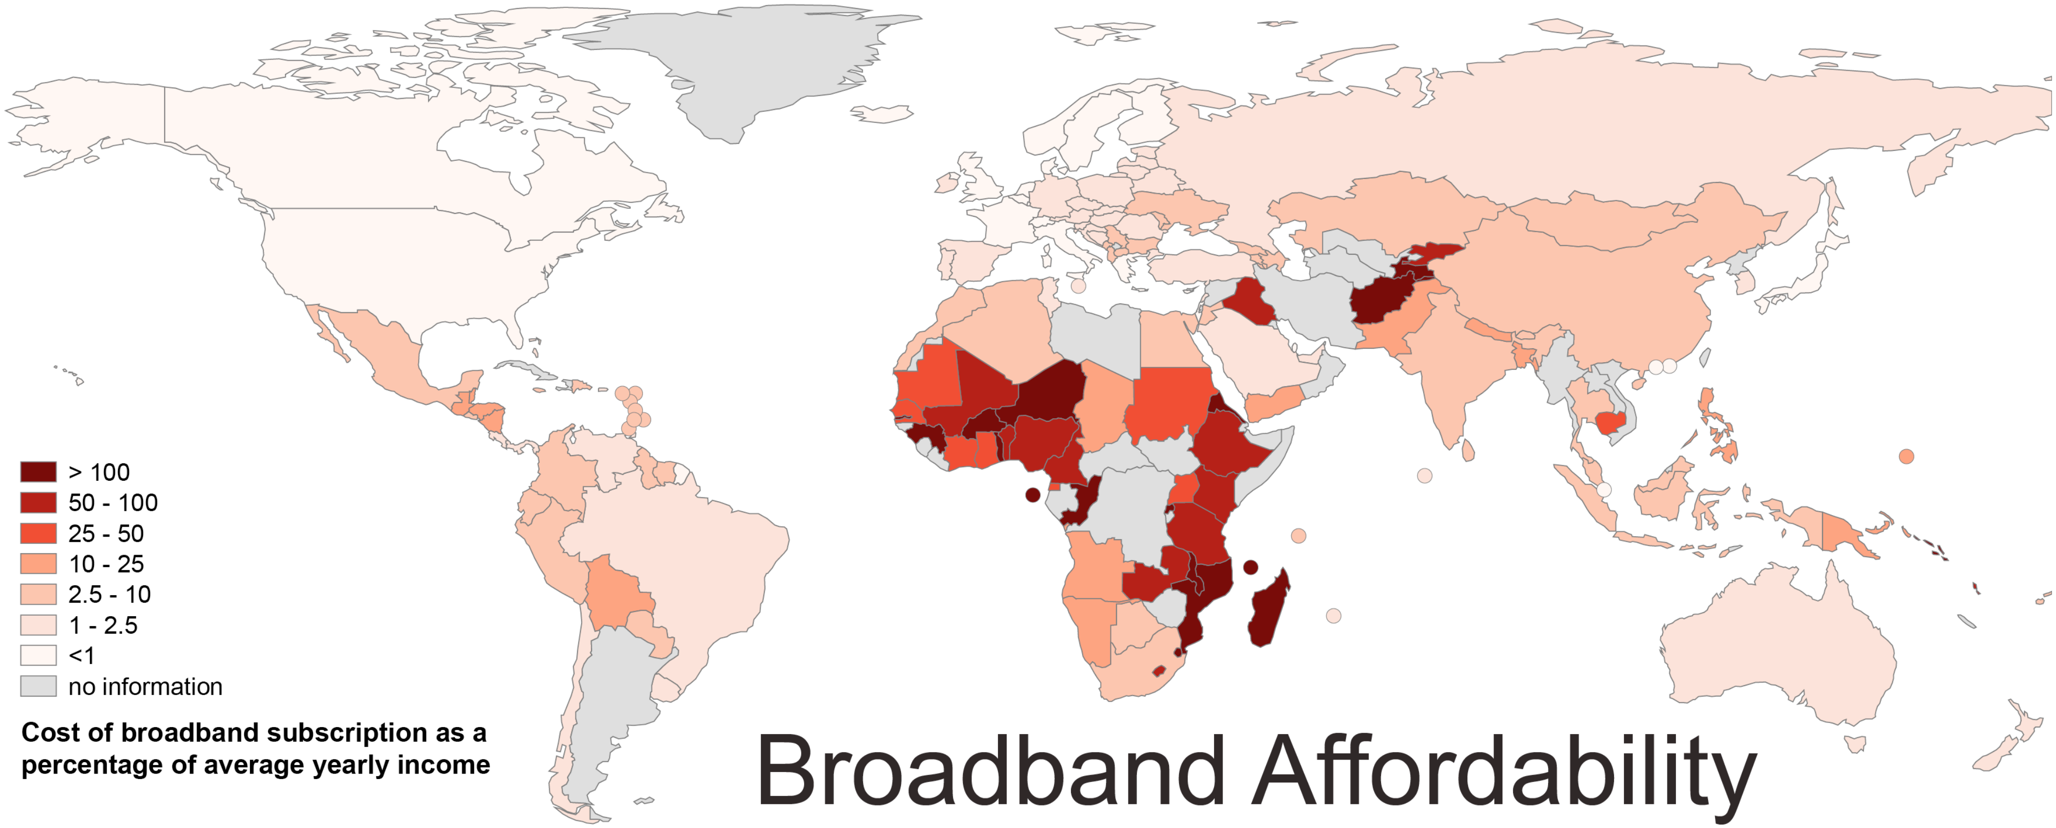

Cost of broadband subscription as a percentage of average yearly income

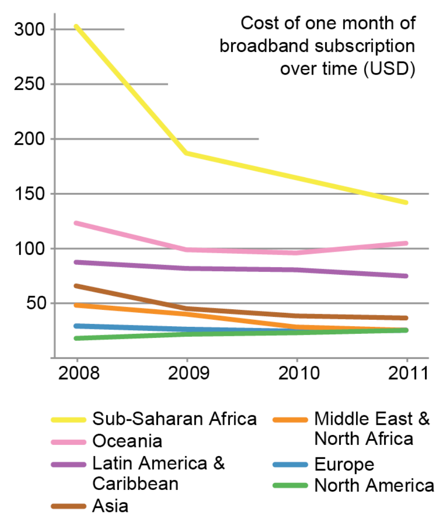

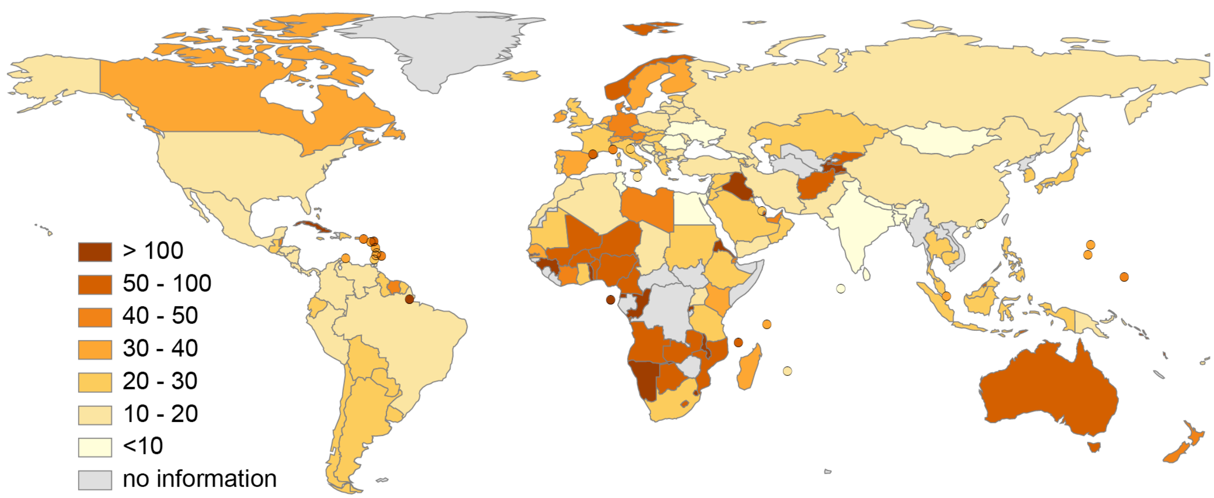

Cost of one month of broadband subscription in 2011 (USD)

Oxford Internet Institute
University of Oxford

by Mark Graham (@geoplace) and Stefano De Sabbata (@maps4thought) Internet Geographies at the Oxford Internet Institute 2014 geography.oii.ox.ac.uk

data sources:
ITU • itu.int
World Bank • data.worldbank.org

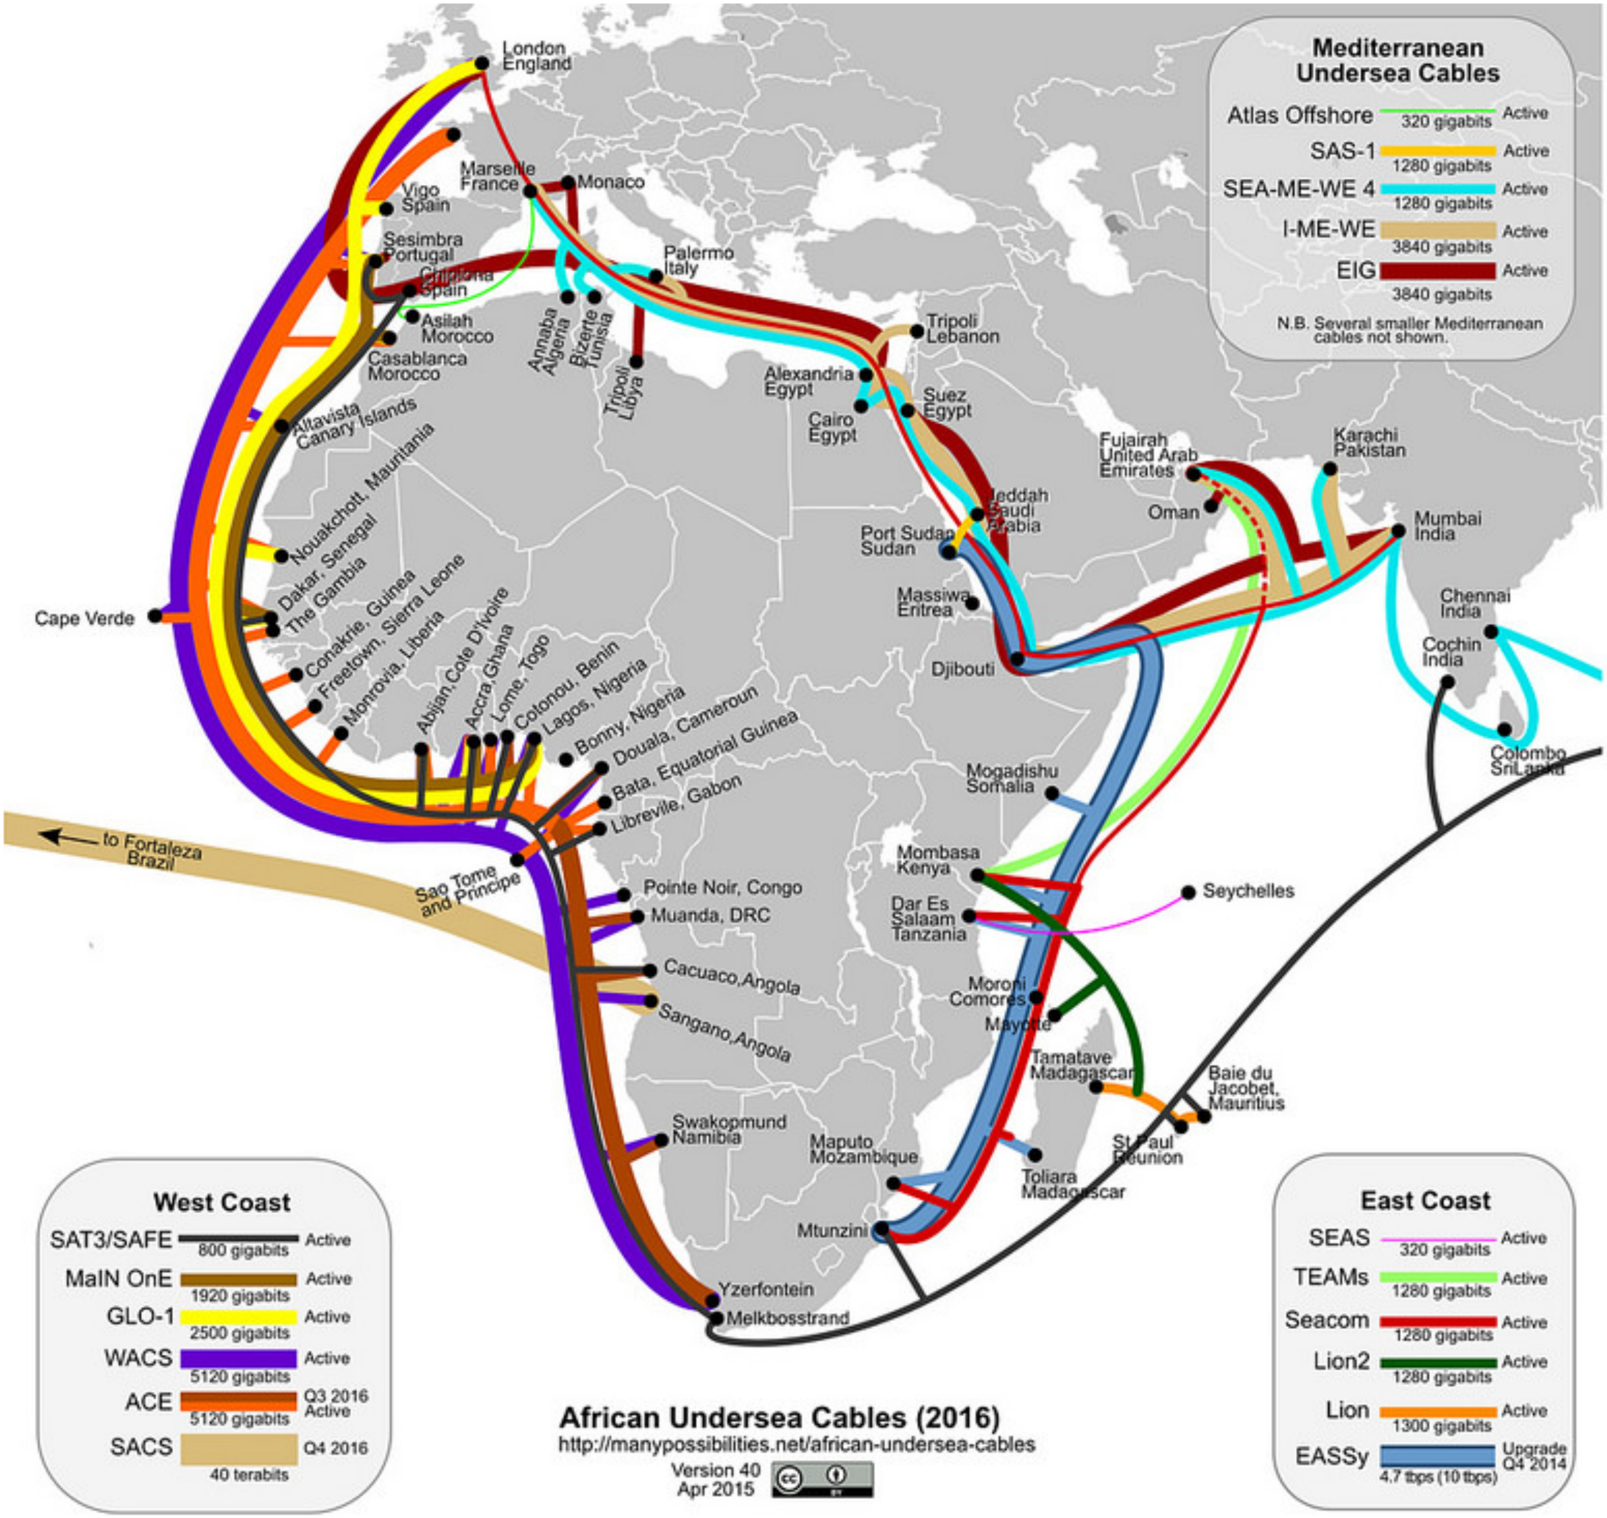
Mediterranean Undersea Cables
Atlas Offshore 320 gigabits Active
SAS-1 1280 gigabits Active
SEA-ME-WE 4 1280 gigabits Active
I-ME-WE 3840 gigabits Active
EIG 3840 gigabits Active
N.B. Several smaller Mediterranean cables not shown.
London England
Marseille France
Monaco
Vigo Spain
Sesimbra Portugal
Chipiona Spain
Asilah Morocco
Casablanca Morocco
Palermo Italy
Annaba Algeria
Bizerte Tunisia
Tripoli Libya
Tripoli Lebanon
Alexandria Egypt
Suez Egypt
Cairo Egypt
Fujairah United Arab Emirates
Karachi Pakistan
Oman
Mumbai India
Jeddah Saudi Arabia
Port Sudan Sudan
Massiwa Eritrea
Djibouti
Chennai India
Cochin India
Colombo SriLanka
Altavista Canary Islands
Nouakchott, Mauritania
Dakar, Senegal
The Gambia
Cape Verde
Conakrie, Guinea
Freetown, Sierra Leone
Monrovia, Liberia
Abijan, Cote D'Ivoire
Accra, Ghana
Lome, Togo
Cotonou, Benin
Lagos, Nigeria
Bonny, Nigeria
Douala, Cameroun
Bata, Equatorial Guinea
Libreville, Gabon
to Fortaleza Brazil
Sao Tome and Principe
Pointe Noir, Congo
Muanda, DRC
Cacuaco, Angola
Sangano, Angola
Swakopmund Namibia
Yzerfontein
Melkbosstrand
Mogadishu Somalia
Mombasa Kenya
Seychelles
Dar Es Salaam Tanzania
Moroni Comores
Mayotte
Tamatave Madagascar
Baie du Jacobet, Mauritius
St Paul Reunion
Toliara Madagascar
Maputo Mozambique
Mtunzini
West Coast
SAT3/SAFE 800 gigabits Active
MaIN OnE 1920 gigabits Active
GLO-1 2500 gigabits Active
WACS 5120 gigabits Active
ACE 5120 gigabits Q3 2016 Active
SACS 40 terabits Q4 2016
East Coast
SEAS 320 gigabits Active
TEAMs 1280 gigabits Active
Seacom 1280 gigabits Active
Lion2 1280 gigabits Active
Lion 1300 gigabits Active
EASSy 4.7 tbps (10 tbps) Upgrade Q4 2014
African Undersea Cables (2016)
http://manypossibilities.net/african-undersea-cables
Version 40
Apr 2015

http://afterfibre.net/

Reliance on Fixed Infrastructure

- Lots of wireless, but only at the very edge
- Lots of *research* into multi-hop wireless
- Little that is practically useful (outside of .mil)

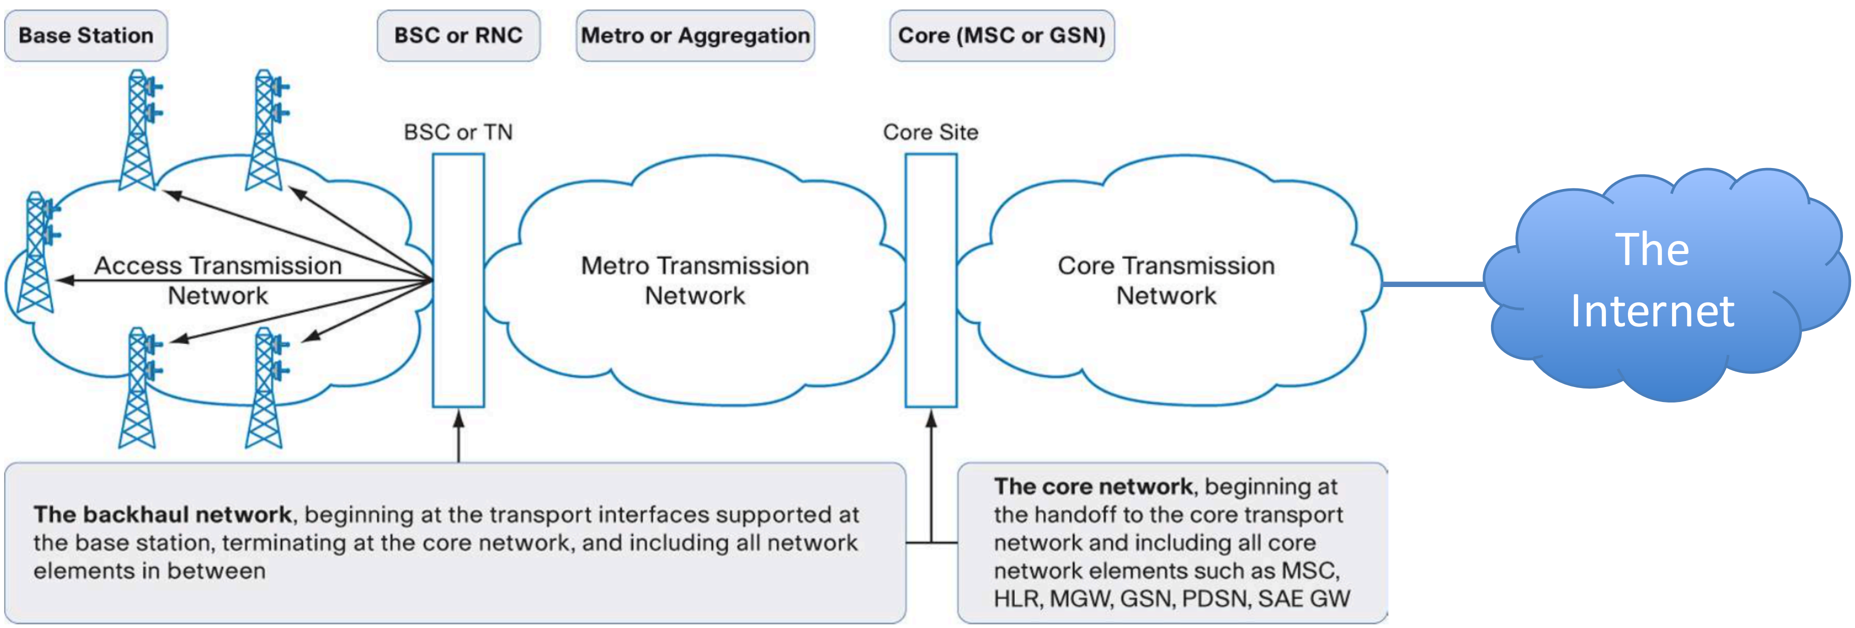

http://www.cisco.com/c/en/us/solutions/collateral/service-provider/unified-ran-backhaul/white_Paper_c11-663732.html

Ad Hoc Connectivity

- Lots of **opportunity** for it
- Was **never easy to use** with/for Internet
  - Addressing
  - Service/information discovery
  - Authentication and confidentiality
- Issues with ad-hoc-style even over fixed infrastructure
  - Peer-to-peer file sharing
- Interplanetary approach? ICN?
  - Services? Content?

Cloudification Doesn't Help

- Clouds are great with good WAN connectivity
  - Easy to deploy & manage services
  - Economy of scale (resources, reliability, etc.)
  - Can do things you can't do easily on prem
- **No clouds with bad WAN connectivity**
- Push clouds everywhere to within a few hops?
- Other ways to enable services/content?
- Critical mass?

AND BEYOND

The Internet Touches All

- **Monitored & programmable physical world**
  - IoT, wearable, mobiles, cars, etc., etc.
  - DNA/RNA-based cell programming
- **All information past & present online, forever**
  - All indexed, searchable, analyzable
  - Better and better machine learning
  - More and more resource capacity & speeds
- **Vast potential for good, vast potential for bad**

Limit Bad, Enable Good – But How?

- Collecting, transmitting, storing data **easy?**
- Keeping data/service online securely **hard**
- Provably deleting data **hard**
- Handling data according to law/policy **hard**
- Specifying law/policy for data/service **hard**
- Agreeing on how to set law/policy **hard**

Final Thought

- Most important Internet research questions in the next 50 years won't be on technology
- **How do we build the Internet to let more individuals attain more of their basic human rights?**
- Not very high on the political agenda
- If anything, Internet is becoming a tool for control

Internet Censorship & Surveillance

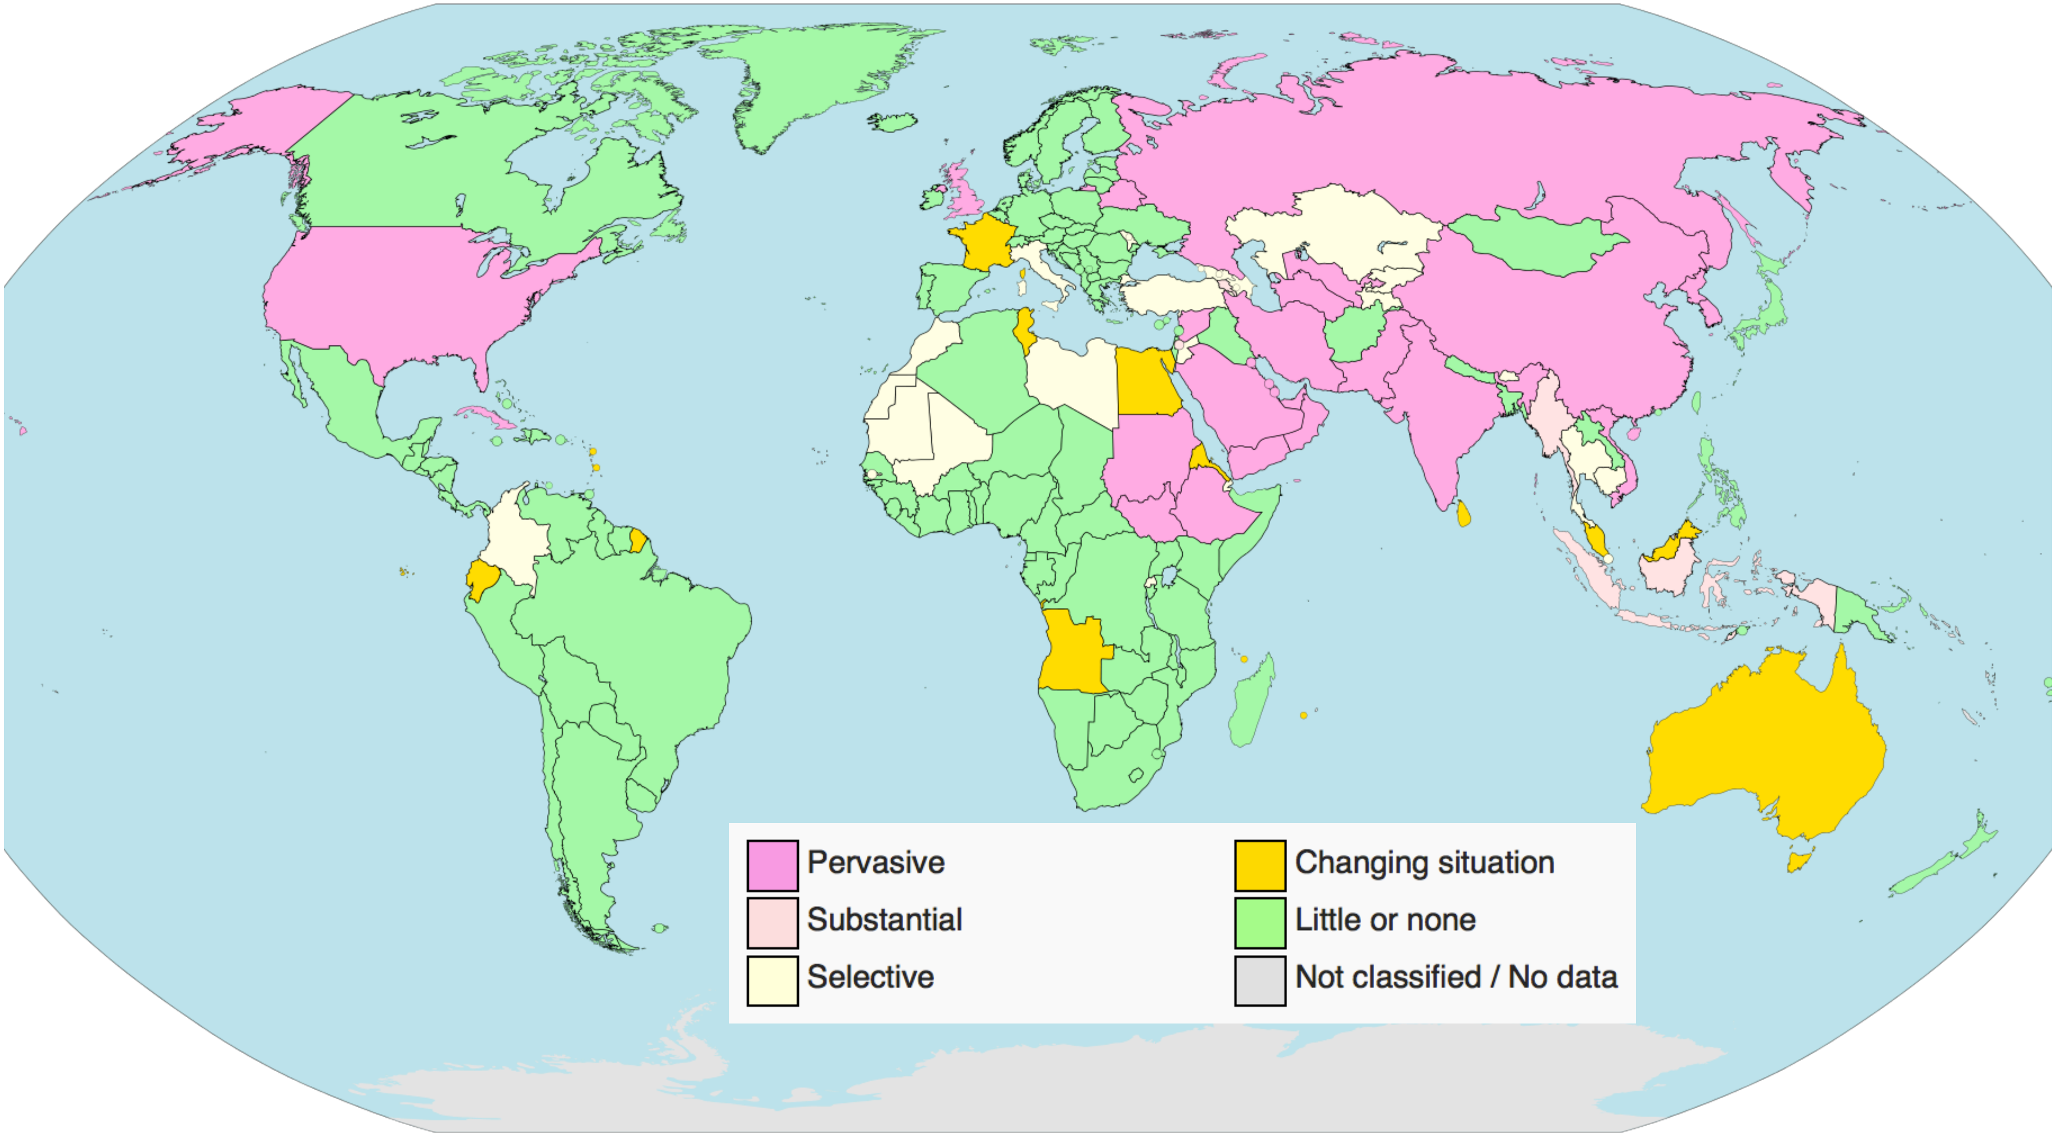

https://commons.wikimedia.org/wiki/File:Internet_Censorship_and_Surveillance_World_Map.svg

lars@netapp.com

THANK YOU