INFO/CS 4302
Web Information Systems

FT 2012

Week 2: Internet History & Architecture

(Part 2)

Theresa Velden

RECAP: INTERNET ARCHITECTURE

Internet Essentials

- IP: common interoperability foundation (addressing and routing)
- Packet switching: flexible, store & forward; difficult to trace
- Architecture layered and modularized
- End-to-end principle: intelligence at the end points
- Domain Name System: two-layer virtual address space
  - IP address (32 bit, 128 bit)
  - Domain name: www.cornell.edu
- Governance by a combo of engineers (IETF), non-governmental agencies, companies and governments

[source: Carl Lagoze, lecture slides INFO4302 Fall 2011]

TCP/IP Protocol Suite

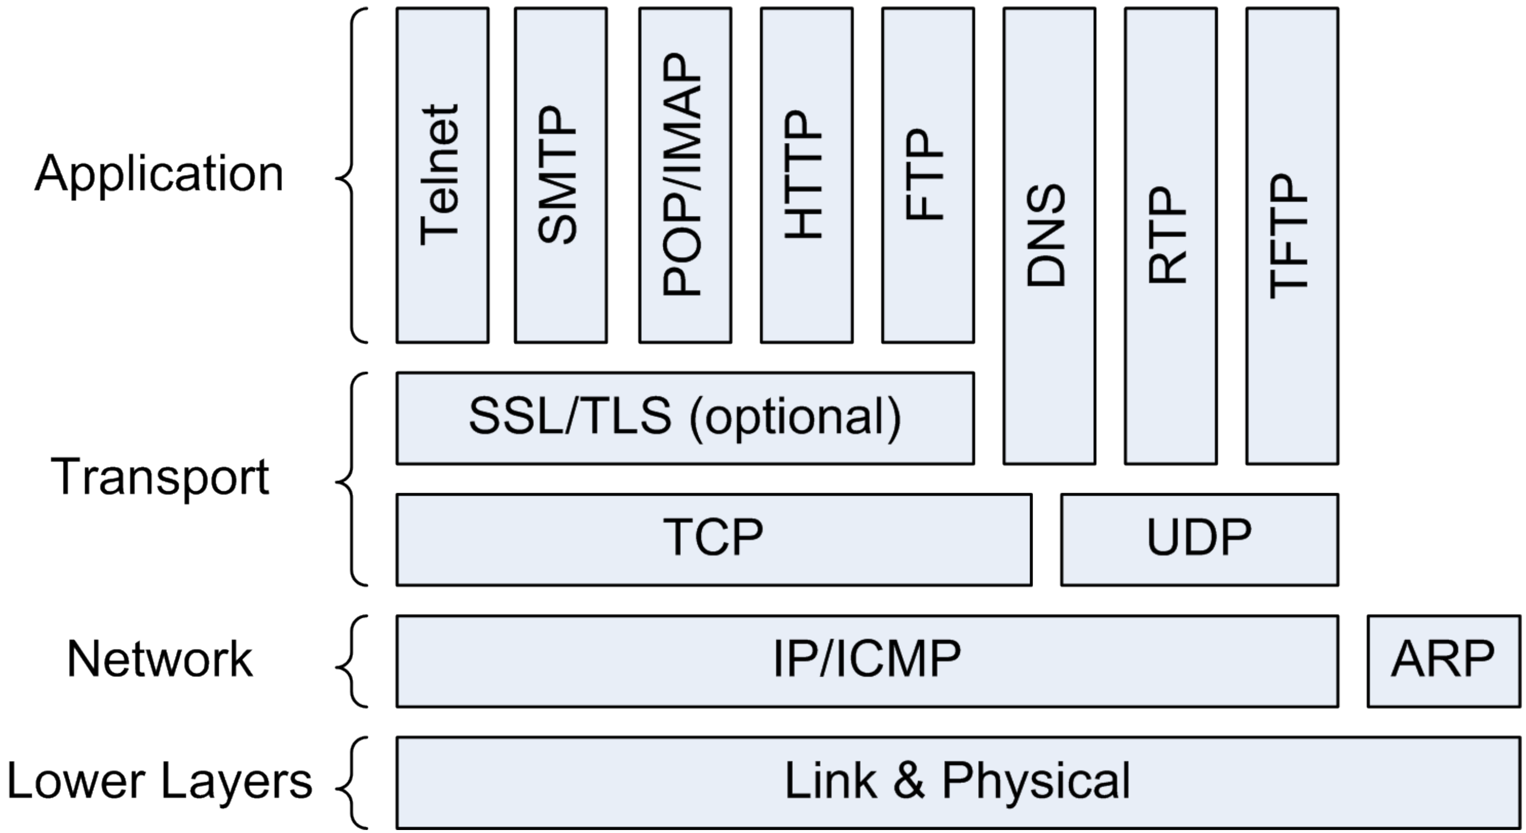

[Source: Eric Wilde & Dolan Mahendran 2011 Creative Commons License]

Well-Known Application Layer Protocols

- Telnet – external terminal interface, RFC 854 (1983)
- FTP – file transfer, RFC 959 (1985)
- SMTP – mail transport, RFC 821 (1982)
- HTTP – distributed, collaborative hypermedia systems, RFC 1945 (1.0 1996), RFC 2616 (1.1 1999)

Network Layer:
ICMP (Internet Control Message Protocol)

ICMP Header

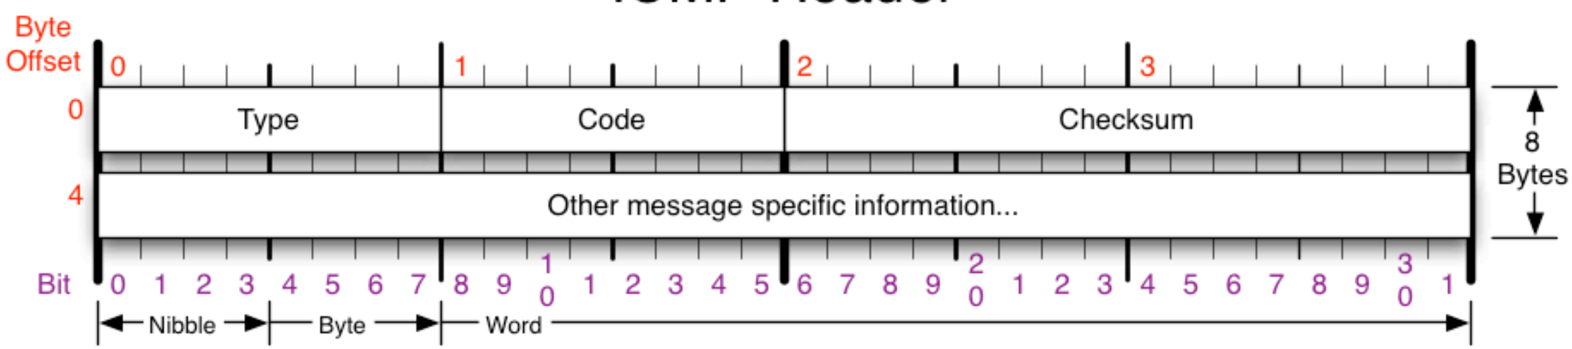

ICMP Message Types

Type Code/Name

0 Echo Reply
3 Destination Unreachable
- 0 Net Unreachable
- 1 Host Unreachable
- 2 Protocol Unreachable
- 3 Port Unreachable
- 4 Fragmentation required, and DF set
- 5 Source Route Failed
- 6 Destination Network Unknown
- 7 Destination Host Unknown
- 8 Source Host Isolated
- 9 Network Administratively Prohibited
- 10 Host Administratively Prohibited
- 11 Network Unreachable for TOS

Type Code/Name

3 Destination Unreachable (continued)
- 12 Host Unreachable for TOS
- 13 Communication Administratively Prohibited

4 Source Quench
5 Redirect
- 0 Redirect Datagram for the Network
- 1 Redirect Datagram for the Host
- 2 Redirect Datagram for the TOS & Network
- 3 Redirect Datagram for the TOS & Host

8 Echo
9 Router Advertisement
10 Router Selection

Type Code/Name

11 Time Exceded
- 0 TTL Exceeded
- 1 Fragment Reassembly Time Exceeded

12 Parameter Problem
- 0 Pointer Problem
- 1 Missing a Required Operand
- 2 Bad Length

13 Timestamp
14 Timestamp Reply
15 Information Request
16 Information Reply
17 Address Mask Request
18 Address Mask Reply
30 Traceroute

Checksum

Checksum of ICMP header

RFC 792

Please refer to RFC 792 for the Internet Control Message protocol (ICMP) specification.

Transport Layer:
User Datagram Protocol

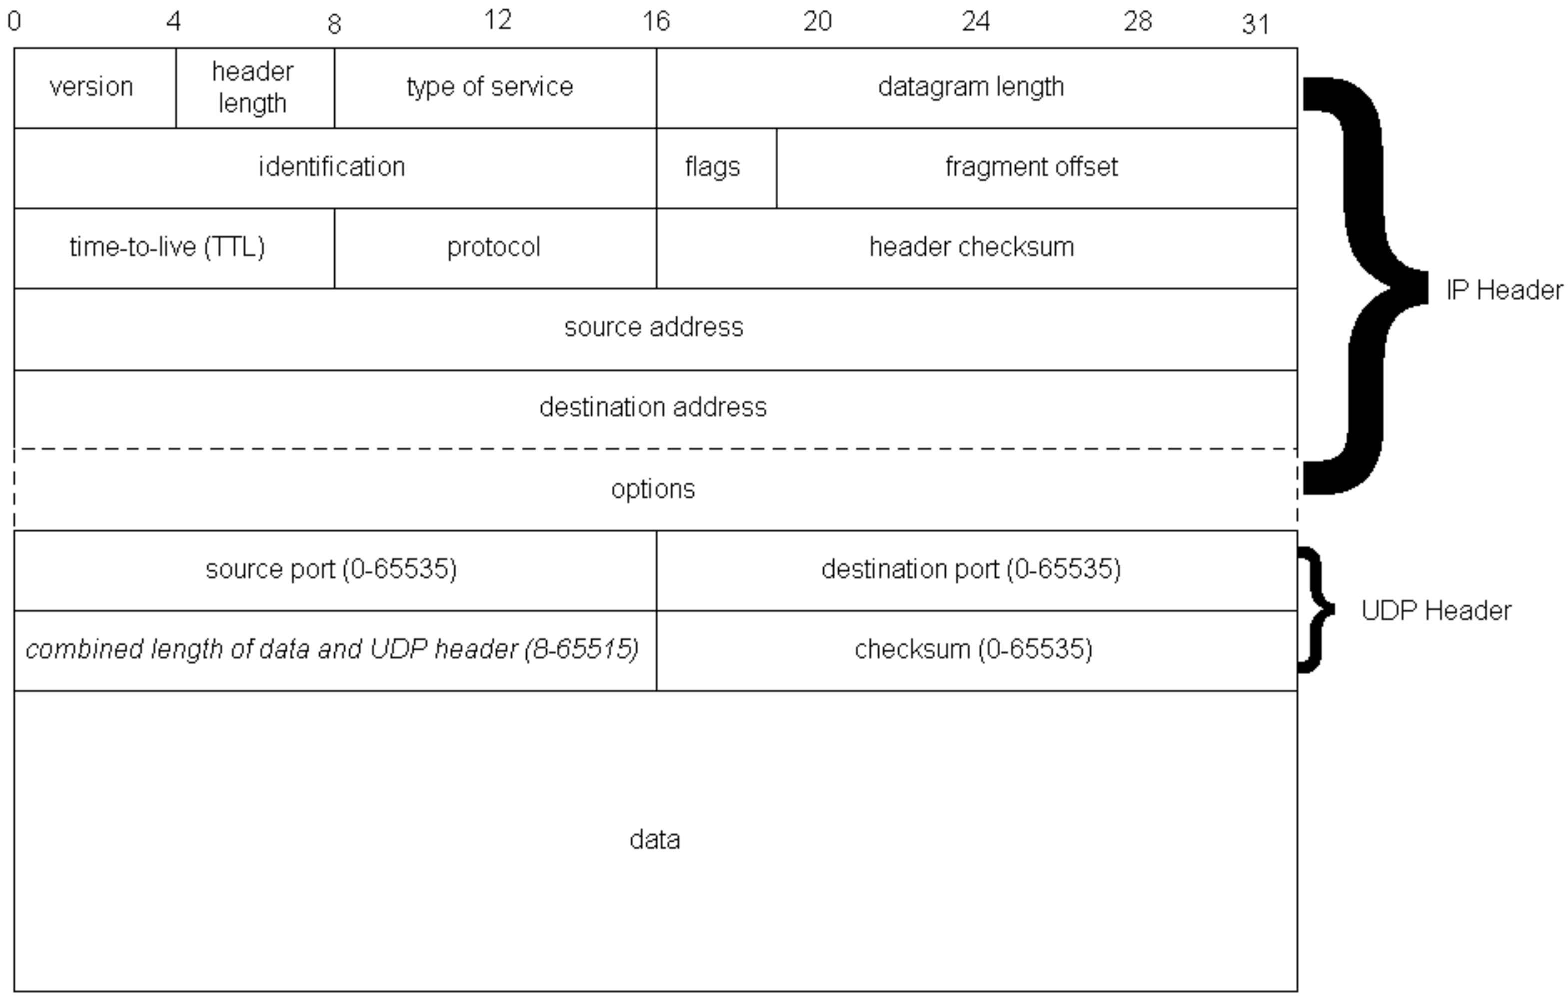

Stateless, unreliable; preferred for streaming

Transport Layer:
TCP Three Way Handshake

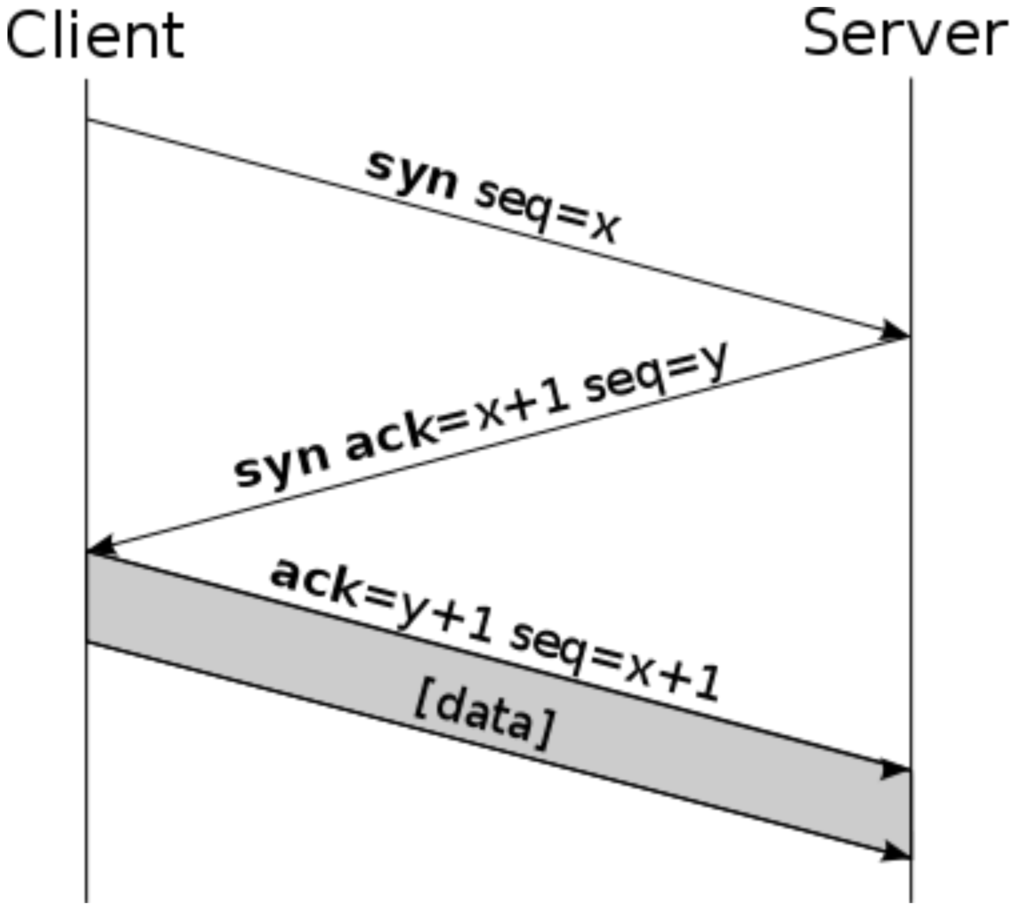

virtual connection, reliable; used for http

Transfer Control Protocol

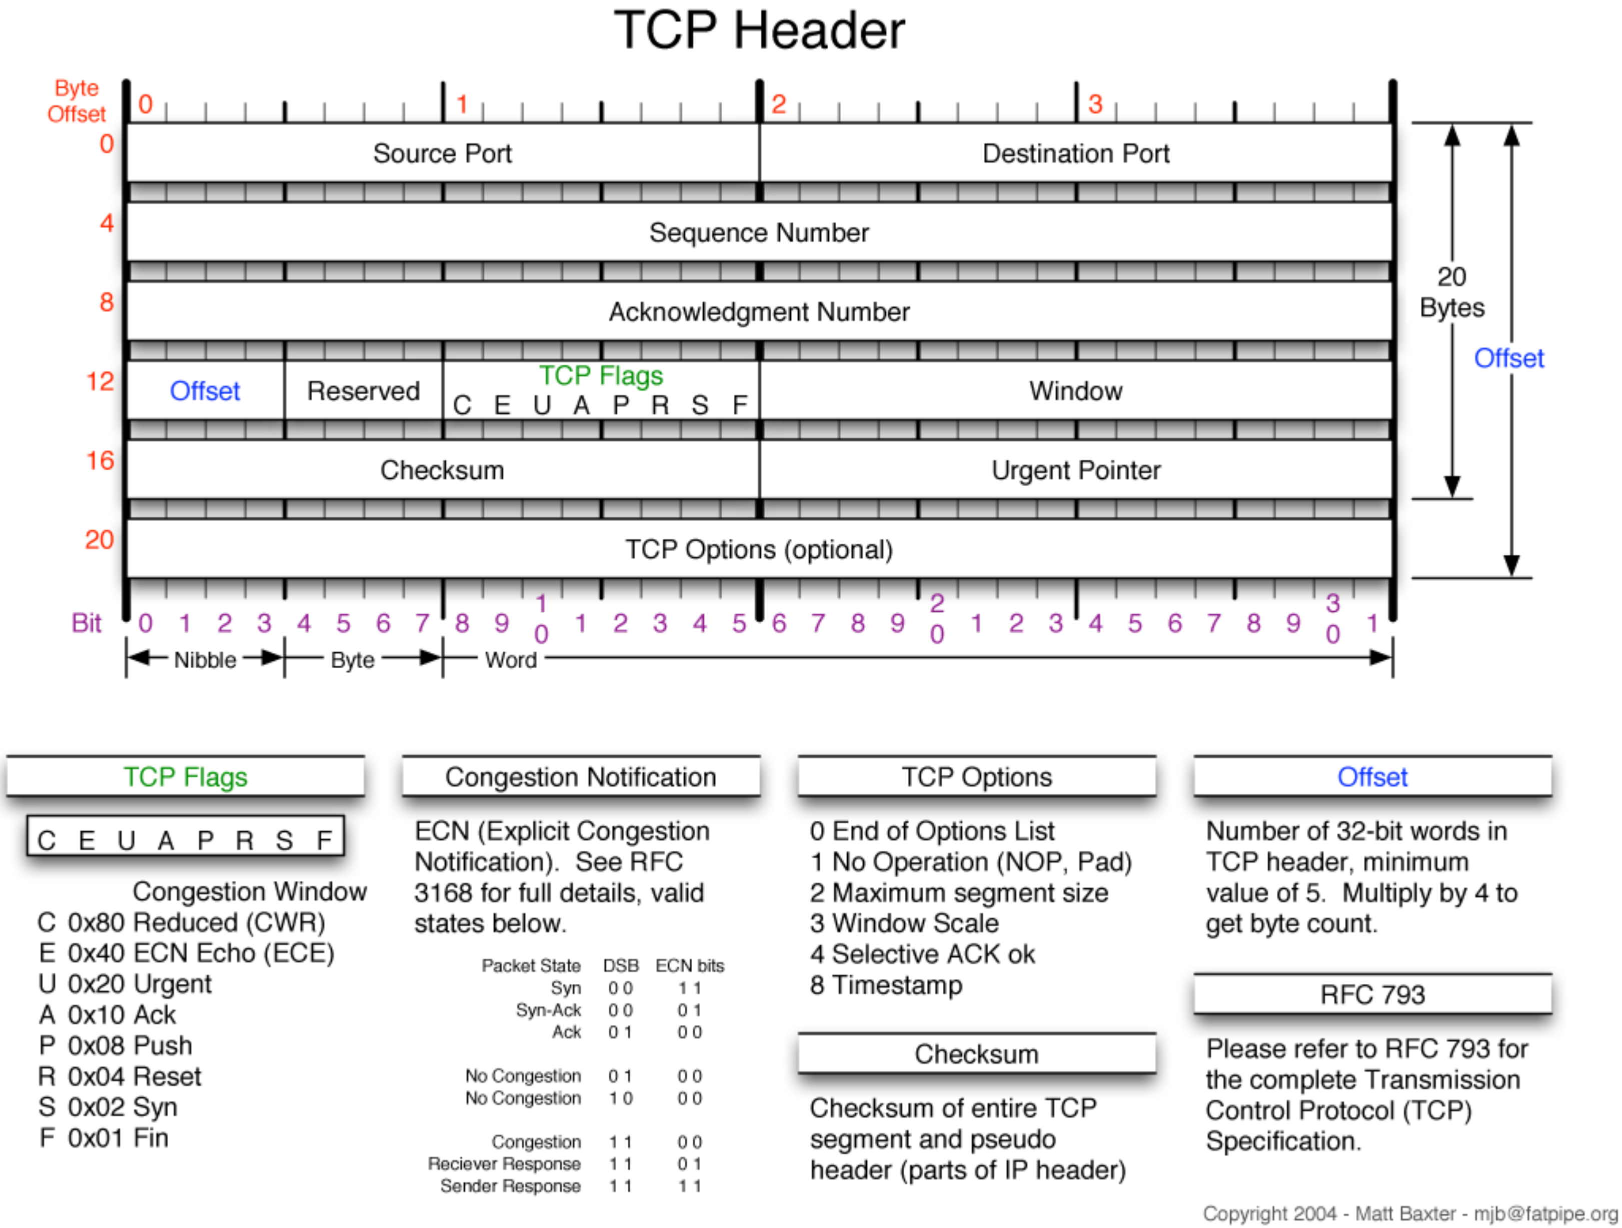

[Source: Eric Wilde & Dolan Mahendran 2011 Creative Commons License]

TCP/UDP Sockets ("Virtual Ports")

- IANA maintains a Service Name and Transport Protocol Port Number Registry:

  "Service names and port numbers are used to distinguish between different services that run over transport protocols such as TCP, UDP, DCCP, and SCTP".

- Services such as: ftp, smtp, whois…
- Needed for binding to applications (IP address/port number, using a protocol such as TCP or UDP)

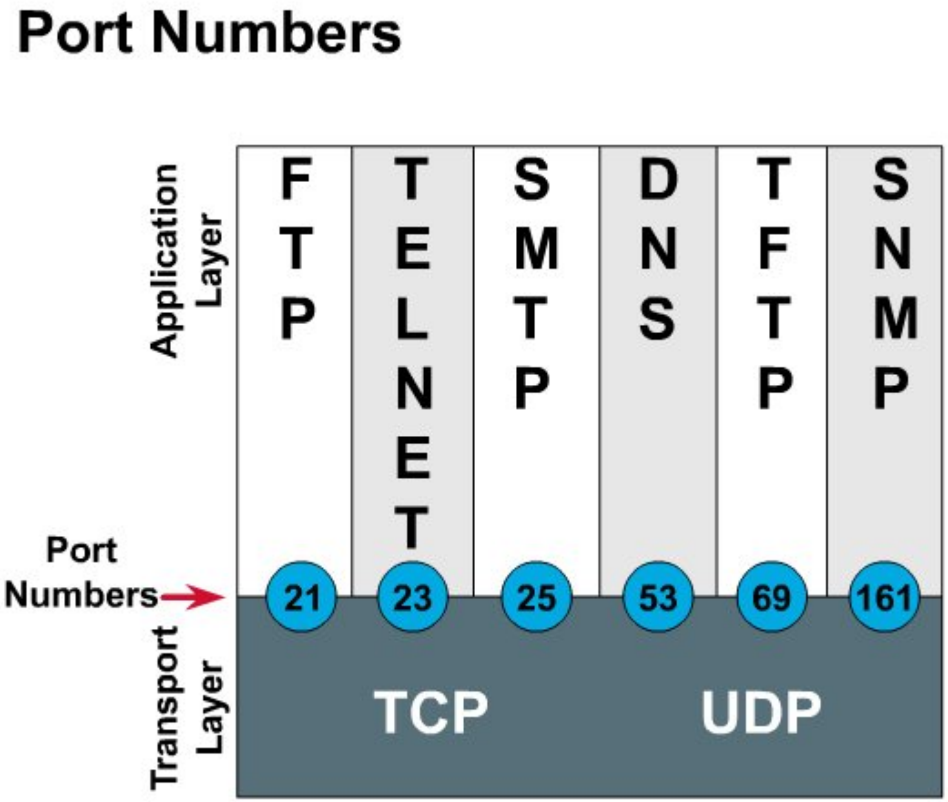

Port Number Listings:

http://www.ietf.org/assignments/service-names-port-numbers/service-names-port-numbers.txt
http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers

INTERNET PROTOCOLS IN ACTION

Which domain name server is my computer using?

```
dhcp103-45:info4302-website theresavelden$ cat /etc/resolv.conf
#
# Mac OS X Notice
#
# This file is not used by the host name and address resolution
# or the DNS query routing mechanisms used by most processes on
# this Mac OS X system.
#
# This file is automatically generated.
#
domain cs.cornell.edu
nameserver 128.84.96.11
nameserver 128.84.86.41
nameserver 128.84.96.10
nameserver 128.84.86.40
```

whois

- **RFC 3912** http://tools.ietf.org/html/rfc3912
- TCP-based transaction-oriented query/response protocol
- Originally: information about registered domain names
- Port Number (IANA):

Service name	Port number	Transport Protocol	Reference
– nicname	43	tcp	Who Is
– nicname	43	udp	Who Is

whois 128.84.86.40?

```
dhcp103-45:info4302-website theresavelden$ whois "n 128.84.86.40"
#
# The following results may also be obtained via:
# http://whois.arin.net/rest/nets;q=128.84.86.40?
showDetails=true&showARIN=false&ext=netref2
#

NetRange:       128.84.0.0 - 128.84.255.255
CIDR:           128.84.0.0/16
OriginAS:
NetName:        CORNELL-NET
NetHandle:      NET-128-84-0-0-1
Parent:         NET-128-0-0-0-0
NetType:        Direct Assignment
RegDate:        1985-05-29
Updated:        2011-10-07
Ref:            http://whois.arin.net/rest/net/NET-128-84-0-0-1

OrgName:        Cornell University
OrgId:          CORNEL
Address:        Cornell Information Technologies
Address:        757 Rhodes Hall
City:           Ithaca
StateProv:      NY
PostalCode:     14853
Country:        US
RegDate:        1984-08-10
Updated:        2011-10-07
Ref:            http://whois.arin.net/rest/org/CORNEL

OrgTechHandle: DEC12-ARIN
OrgTechName:   Eckstrom, Daniel
OrgTechPhone:  +1-607-255-5902
OrgTechEmail:  de10@cornell.edu
OrgTechRef:    http://whois.arin.net/rest/poc/DEC12-ARIN
```

```
dhcp103-45:info4302-website theresavelden$ whois "n 128.84.86.40"
#
# The following results may also be obtained via:
# http://whois.arin.net/rest/nets;q=128.84.86.40?
showDetails=true&showARIN=false&ext=netref2
#

NetRange:       128.84.0.0 - 128.84.255.255
CIDR:           128.84.0.0/16
OriginAS:
NetName:        CORNELL-NET
NetHandle:      NET-128-84-0-0-1
Parent:         NET-128-0-0-0-0
NetType:        Direct Assignment
RegDate:        1985-05-29
Updated:        2011-10-07
Ref:            http://whois.arin.ne

OrgName:        Cornell University
OrgId:          CORNEL
Address:        Cornell Information
Address:        757 Rhodes Hall
City:           Ithaca
StateProv:      NY
PostalCode:     14853
Country:        US
RegDate:        1984-08-10
Updated:        2011-10-07
Ref:            http://whois.arin.ne

OrgTechHandle: DEC12-ARIN
OrgTechName:   Eckstrom, Daniel
OrgTechPhone:  +1-607-255-5902
OrgTechEmail:  de10@cornell.edu
OrgTechRef:    http://whois.arin.net/rest/poc/DEC12-ARIN
```

Does Cornell have to have a domain name server?

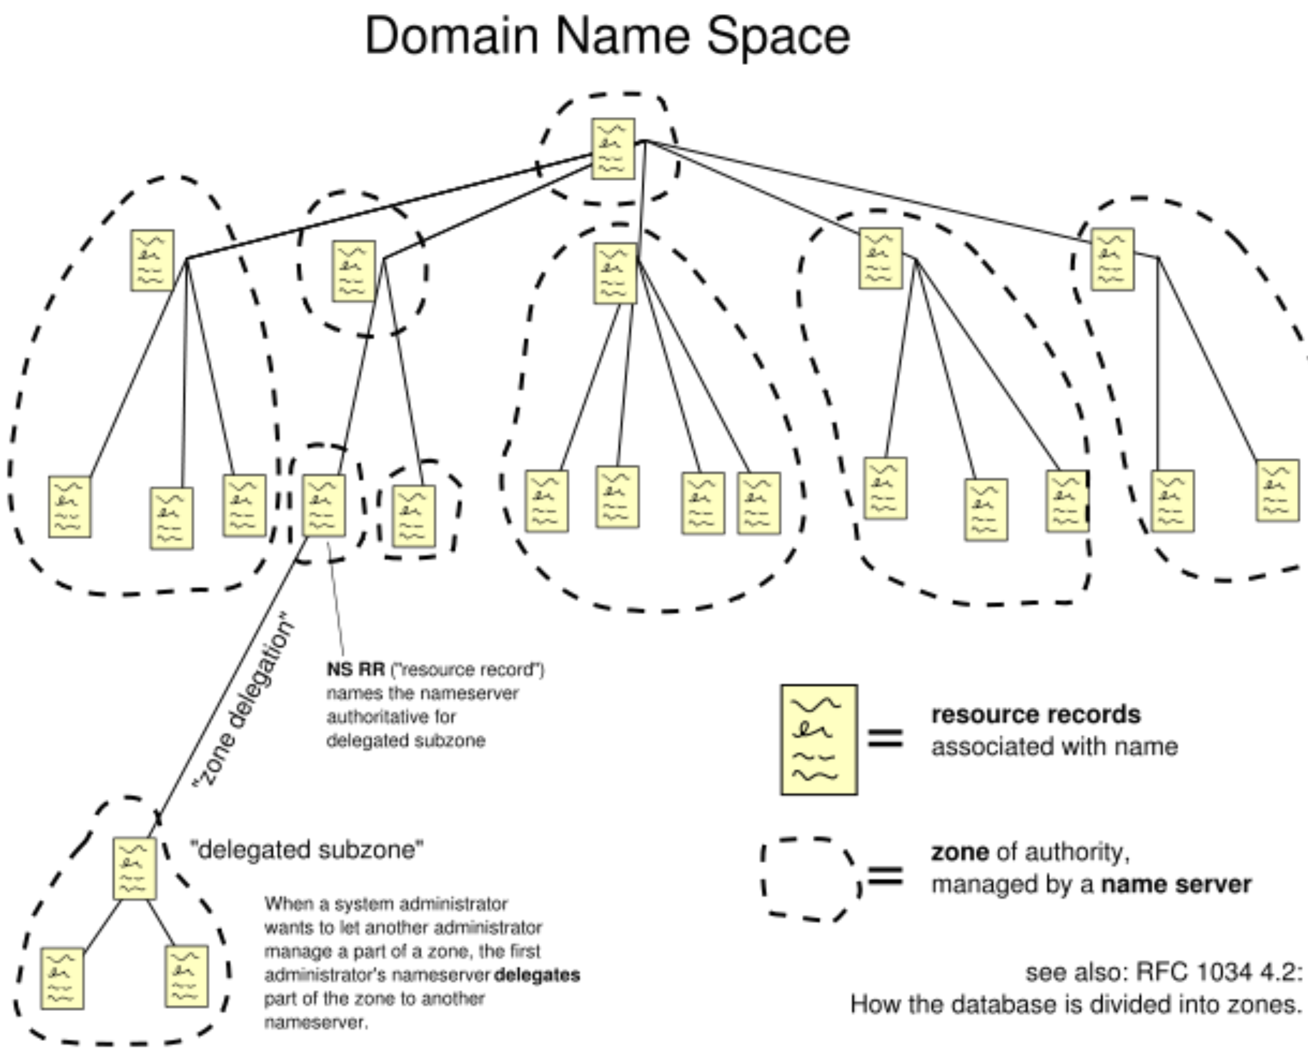

How is my DNS lookup routed?

- Choose domain name server to send request to
- Public domain name server IP address = **8.8.8.8**   (provided courtesy of google)
- Domain name I want to resolve: **cs.cornell.edu**
- DNS lookup utility: “dig”

Tracing Domain Name Resolution (1)

dig +trace @8.8.8.8 cs.cornell.edu

domain name

public name server by google

```
dhcp103-45:info4302-website theresavelden$ dig +trace @8.8.8.8 cs.cornell.edu

; <<>> DiG 9.6-ESV-R4-P3 <<>> +trace @8.8.8.8 cs.cornell.edu
; (1 server found)
;; global options: +cmd
.            11187   IN   NS   h.root-servers.net.
.            11187   IN   NS   g.root-servers.net.
.            11187   IN   NS   a.root-servers.net.
.            11187   IN   NS   k.root-servers.net.
.            11187   IN   NS   m.root-servers.net.
.            11187   IN   NS   e.root-servers.net.
.            11187   IN   NS   d.root-servers.net.
.            11187   IN   NS   j.root-servers.net.
.            11187   IN   NS   f.root-servers.net.
.            11187   IN   NS   l.root-servers.net.
.            11187   IN   NS   i.root-servers.net.
.            11187   IN   NS   b.root-servers.net.
.            11187   IN   NS   c.root-servers.net.
;; Received 228 bytes from 8.8.8.8#53(8.8.8.8) in 20 ms
```

from public name server

Tracing Domain Name Resolution (2)

```
edu.            172800   IN   NS   g.edu-servers.net.
edu.            172800   IN   NS   c.edu-servers.net.
edu.            172800   IN   NS   l.edu-servers.net.
edu.            172800   IN   NS   f.edu-servers.net.
edu.            172800   IN   NS   d.edu-servers.net.
edu.            172800   IN   NS   a.edu-servers.net.
;; Received 267 bytes from 192.36.148.17#53(i.root-servers.net) in
113 ms
```

from root server

```
cornell.edu.      172800   IN   NS   bigred.cit.cornell.edu.
cornell.edu.      172800   IN   NS   dns.cit.cornell.edu.
cornell.edu.      172800   IN   NS   drdns.cit.cornell.edu.
;; Received 143 bytes from 192.31.80.30#53(d.edu-servers.net) in 25
ms
```

from top level domain server

```
cs.cornell.edu.       86400    IN   A    128.84.96.11
cs.cornell.edu.       86400    IN   A    128.84.96.12
cs.cornell.edu.       86400    IN   A    128.84.96.13
cs.cornell.edu.       86400    IN   A    128.84.96.10
cs.cornell.edu.       432000   IN   NS   bigred.cit.cornell.edu.
cs.cornell.edu.       432000   IN   NS   dns.cit.cornell.edu.
cs.cornell.edu.       432000   IN   NS   drdns.cit.cornell.edu.
cs.cornell.edu.       432000   IN   NS   cudns.cit.cornell.edu.
;; Received 243 bytes from 192.35.82.53#53(dns.cit.cornell.edu) in 1 ms
```

Tracing Domain Name Resolution (2)

```
edu.            172800  IN  NS  g.edu-servers.net.
edu.            172800  IN  NS  c.edu-servers.net.
edu.            172800  IN  NS  l.edu-servers.net.
edu.            172800  IN  NS  f.edu-servers.net.
edu.            172800  IN  NS  d.edu-servers.net.
edu.            172800  IN  NS  a.edu-servers.net.
;; Received 267 bytes from 192.36.148.17#53(i.root-servers.net) in 113 ms

cornell.edu.    172800  IN  NS  bigred.cit.cornell.edu.
cornell.edu.    172800  IN  NS  dns.cit.cornell.edu.
cornell.edu.    172800  IN  NS  drdns.cit.cornell.edu.
;; Received 143 bytes from 192.31.80.30#53(d.edu-servers.net) in 25 ms

cs.cornell.edu.     86400   IN  A   128.84.96.11
cs.cornell.edu.     86400   IN  A   128.84.96.12
cs.cornell.edu.     86400   IN  A   128.84.96.13
cs.cornell.edu.     86400   IN  A   128.84.96.10
cs.cornell.edu.     432000  IN  NS  bigred.cit.cornell.edu
cs.cornell.edu.     432000  IN  NS  dns.cit.cornell.edu.
cs.cornell.edu.     432000  IN  NS  drdns.cit.cornell.edu.
cs.cornell.edu.     432000  IN  NS  cudns.cit.cornell.edu.
;; Received 243 bytes from 192.35.82.53#53(dns.cit.cornell.edu) in 1 ms
```

#53 ???

TCP/UDP Sockets ("Virtual Ports")

19	TCP	UDP	Character Generator Protocol (CHARGEN)
20	TCP		FTP—data transfer
21	TCP		FTP—control (command)
22	TCP		Secure Shell (SSH)—used for secure logins, file transfers (scp, sftp) and port forwarding
23	TCP		Telnet protocol—unencrypted text communications
24	TCP	UDP	Priv-mail : any private mail system.
25	TCP		Simple Mail Transfer Protocol (SMTP)—used for e-mail routing between mail servers
26	TCP	UDP	Unassigned
27	TCP	UDP	NSW User System FE
29	TCP	UDP	MSG ICP
33	TCP	UDP	Display Support Protocol
35	TCP	UDP	Any private printer server protocol
37	TCP	UDP	TIME protocol
39	TCP	UDP	Resource Location Protocol[7] (RLP)—used for determining the location of higher level services
40	TCP	UDP	Unassigned
42	TCP	UDP	ARPA Host Name Server Protocol
42	TCP	UDP	Windows Internet Name Service
43	TCP		WHOIS protocol
47	TCP	UDP	NI FTP[7]
49	TCP	UDP	TACACS Login Host protocol
50	TCP	UDP	Remote Mail Checking Protocol[8]
51	TCP	UDP	IMP Logical Address Maintenance
52	TCP	UDP	XNS (Xerox Network Systems) Time Protocol
53	TCP	UDP	Domain Name System (DNS)
54	TCP	UDP	XNS (Xerox Network Systems) Clearinghouse

http://en.wikipedia.org/wiki/List_of_TCP_and_UDP_port_numbers

What is the Internet route from cornell.edu to destinations abroad?

Examples:

- Max-Planck-Society (Germany)
- Chinese Academy of Sciences? (China)

How traceroute works:

- Sequence of packets send to a destination server
  - Windows: ICMP(Internet Control Message Protocol) echo requests
  - UNIX/LINUX: UDP + (unlikely) port number
- Count down of TTL (time-to-live) with each hop (port number in UNIX/LINUX case)
- When TTL=0, ICMP error message ‘TTL exceeded in transit’ is returned
- Traceroute determines IP of last hop from source field of error message, and time stamp for round trip
- Set of successive packets, each with TTL+1
- Latency:
  - Per round trip (return may differ from outgoing path and speed)
  - Physical distance matter! light through fiber = ~200,000km/sec, → 200km (or 125 miles) per millisecond

traceroute

```
dhcp103-45:info4302-website theresavelden$ traceroute -I mpg.de
traceroute to mpg.de (134.76.31.209), 64 hops max, 52 byte packets
 1  csgate8 (128.84.103.2)  1.842 ms  1.600 ms  1.536 ms
 2  aerberus-rt (128.84.103.146)  0.883 ms  0.697 ms  0.753 ms
 3  rhodes1-6500-vl339.net.cornell.edu (128.84.154.1)  1.037 ms  0.918 ms
 4  core1-6500-te3-3.net.cornell.edu (128.253.222.161)  1.095 ms  1.256 ms
 5  nyc1-6500-vl11.net.cornell.edu (128.253.222.14)  9.065 ms  8.945 ms  9
 6 newy-nlr-vl4000.nelr.net (192.35.82.129)  8.998 ms  8.901 ms  8.835 ms
 7 216.24.184.86 (216.24.184.86)  83.941 ms  84.049 ms  84.031 ms
 8  so-2-0-0.rt1.fra.de.geant2.net (62.40.112.9)  90.824 ms  90.959 ms  10
 9  dfn-gw.rt1.fra.de.geant.net (62.40.124.34)  91.677 ms  92.054 ms  91.6
 10  zr-han1-te0-0-0-1.x-win.dfn.de (188.1.145.214)  103.041 ms  103.042 m
 11  xr-goe1-te1-1.x-win.dfn.de (188.1.144.70)  104.882 ms  104.831 ms  10
 12  kr-gwdg.x-win.dfn.de (188.1.231.126)  104.841 ms  104.817 ms  104.901
 13  gr-gwdg1-fvrf-int.gwdg.de (134.76.249.185)  104.849 ms  104.913 ms  1
 14  public.web.mpg.de (134.76.31.209)  117.142 ms  116.965 ms  117.271 ms
```

traceroute

```
dhcp103-45:info4302-website theresavelden$ traceroute -I mpg.de
traceroute to mpg.de (134.76.31.209), 64 hops max, 52 byte packets
 ***************** CORNELL LAN ************************
 1  csgate8 (128.84.103.2)  1.842 ms  1.600 ms  1.536 ms
 2  aerberus-rt (128.84.103.146)  0.883 ms  0.697 ms  0.753 ms
 3  rhodes1-6500-vl339.net.cornell.edu (128.84.154.1)  1.037 ms  0.918 ms
 4  core1-6500-te3-3.net.cornell.edu (128.253.222.161)  1.095 ms  1.256 ms
 5  nyc1-6500-vl11.net.cornell.edu (128.253.222.14)  9.065 ms  8.945 ms  9
 ***************** CORNELL DMZ ********************
 6 newy-nlr-vl4000.nelr.net (192.35.82.129)  8.998 ms  8.901 ms  8.835 ms
 ***************** National LambdaRail ********************
 7 216.24.184.86 (216.24.184.86)  83.941 ms  84.049 ms  84.031 ms
 ***************** GEANT European Backbone **********************
 8  so-2-0-0.rt1.fra.de.geant2.net (62.40.112.9)  90.824 ms  90.959 ms  10
 9  dfn-gw.rt1.fra.de.geant.net (62.40.124.34)  91.677 ms  92.054 ms  91.6
 ***************** German Research Network X-Win **********************
 10  zr-han1-te0-0-0-1.x-win.dfn.de (188.1.145.214)  103.041 ms  103.042 ms
 11  xr-goe1-te1-1.x-win.dfn.de (188.1.144.70)  104.882 ms  104.831 ms  10
 12  kr-gwdg.x-win.dfn.de (188.1.231.126)  104.841 ms  104.817 ms  104.901
 ***************** Max Planck Computing Center ********************
 13  gr-gwdg1-fvrf-int.gwdg.de (134.76.249.185)  104.849 ms  104.913 ms  10
 14  public.web.mpg.de (134.76.31.209)  117.142 ms  116.965 ms  117.271 ms
```

Internet Routing Architeture

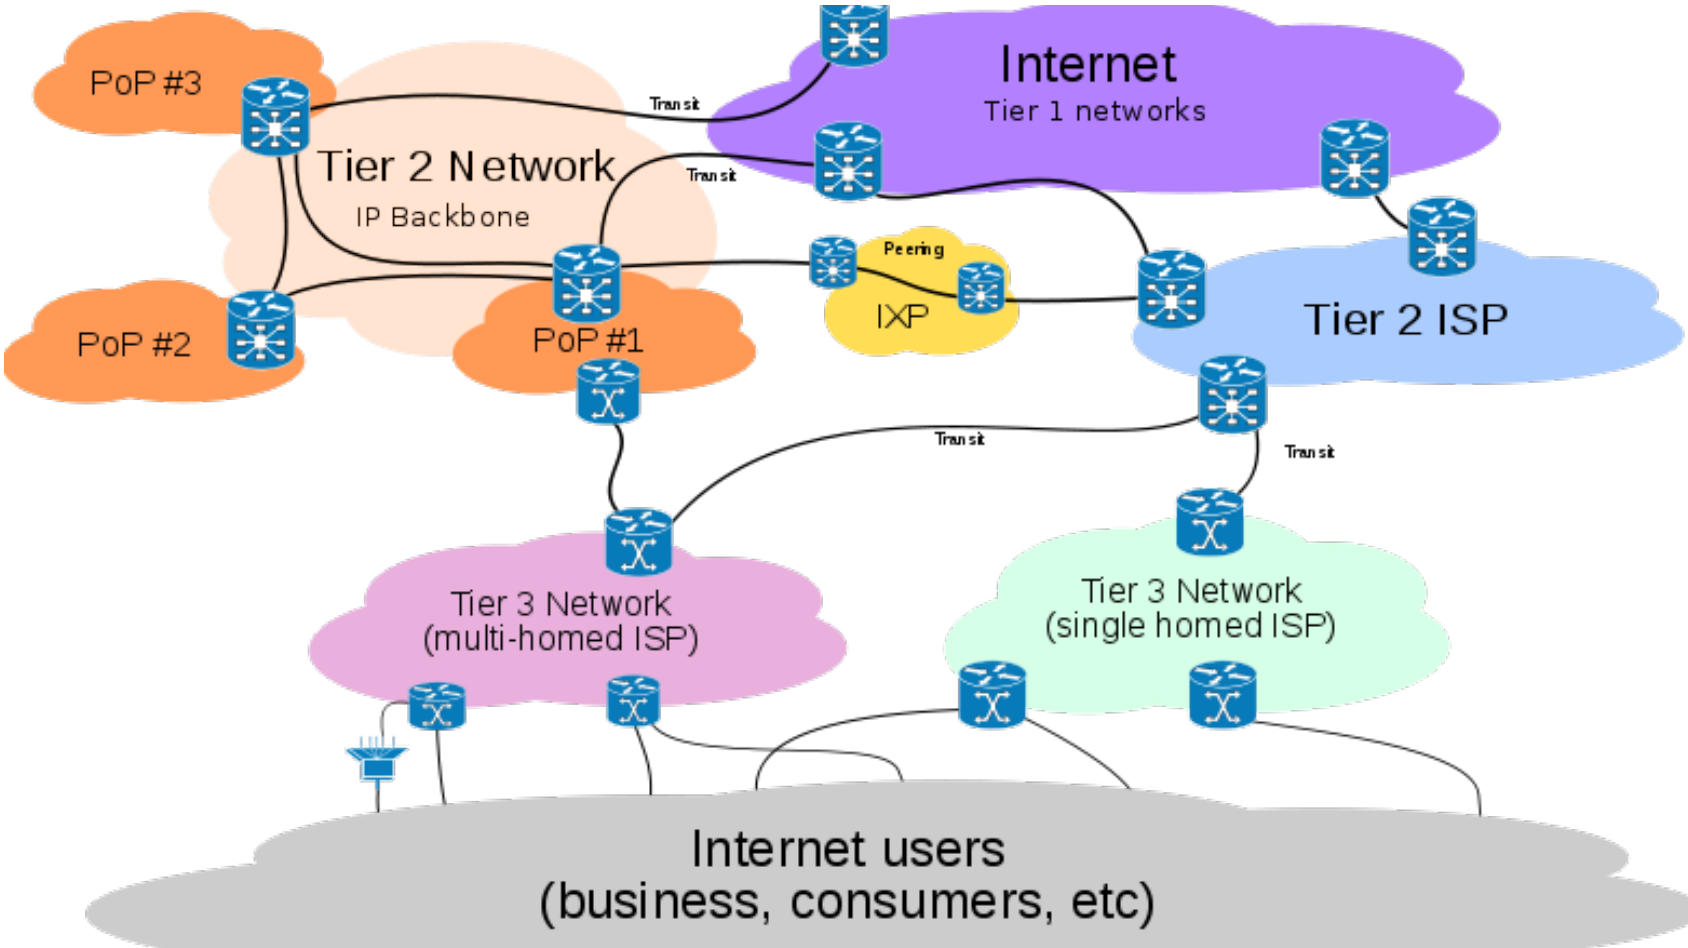

Source & further reading: http://en.wikipedia.org/wiki/Tier_1_network

traceroute Visualization

Traceroute visualization

cs.cornell – www.mpg.de

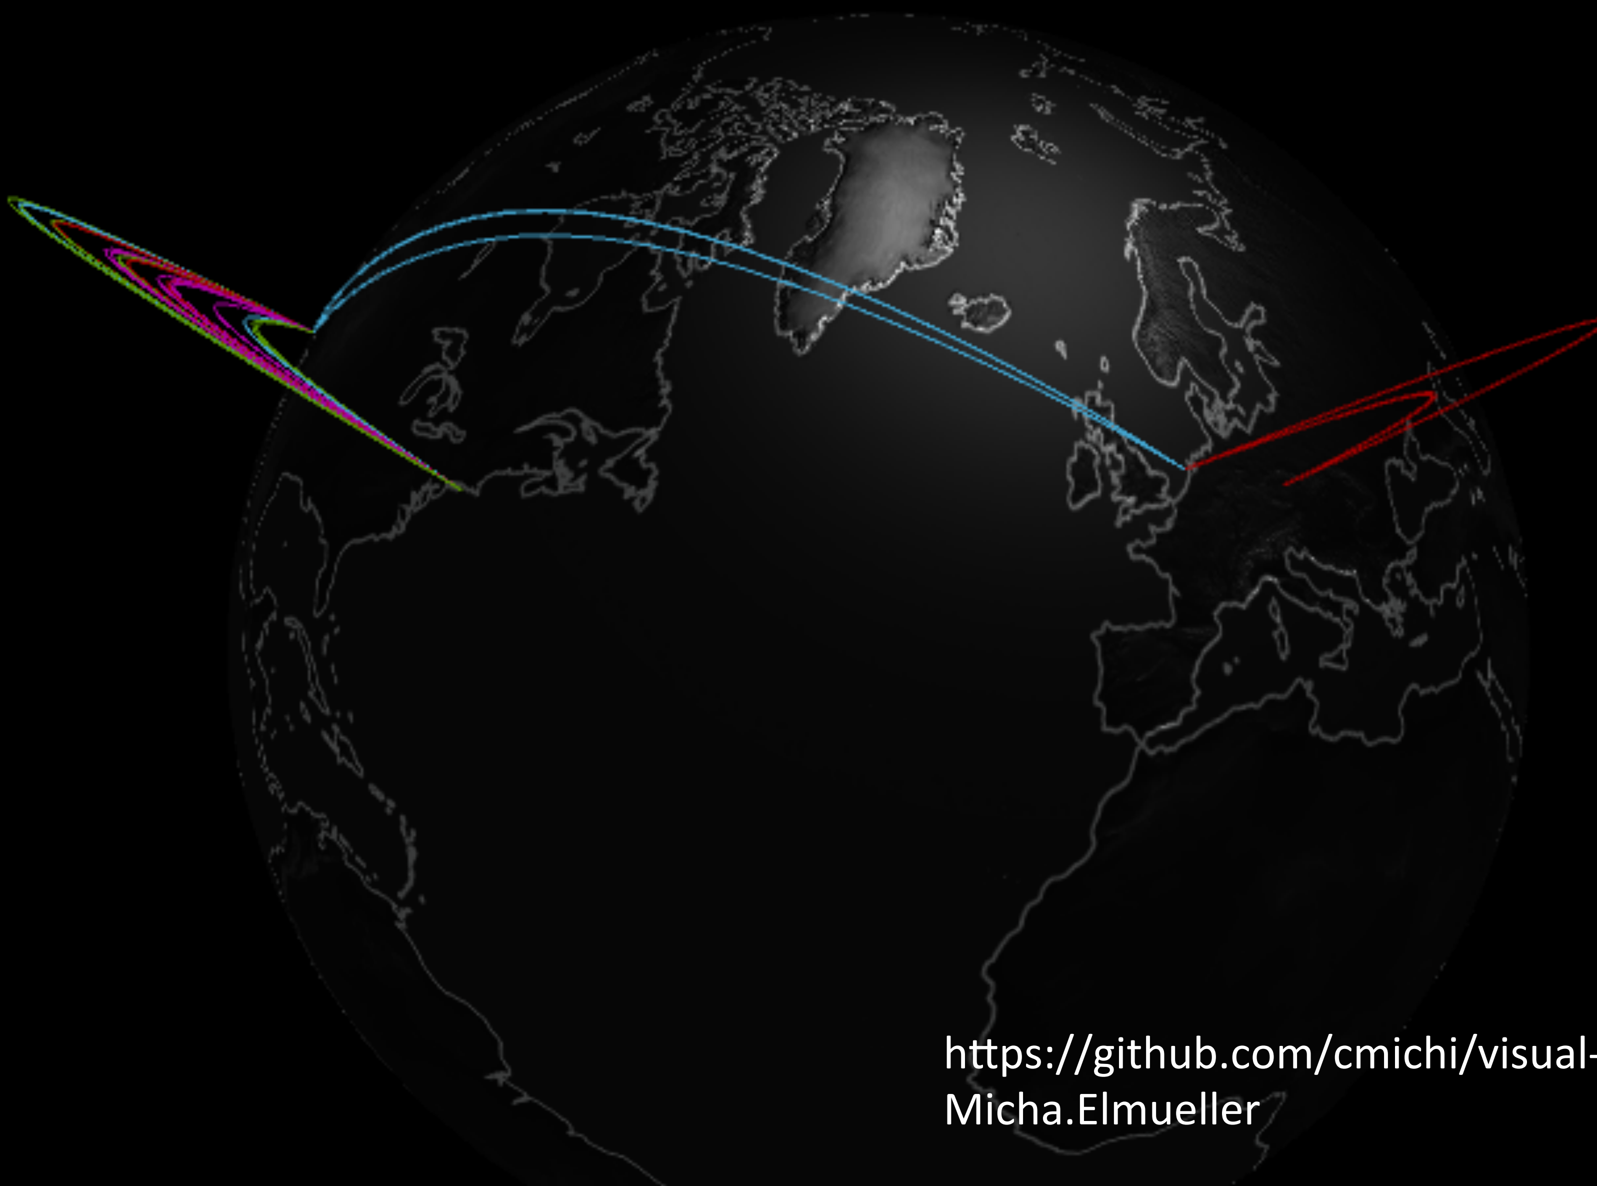

https://github.com/cmichi/visual-traceroute
Micha.Elmueller

Traceroute visualization

cs.cornell – www.cas.cn

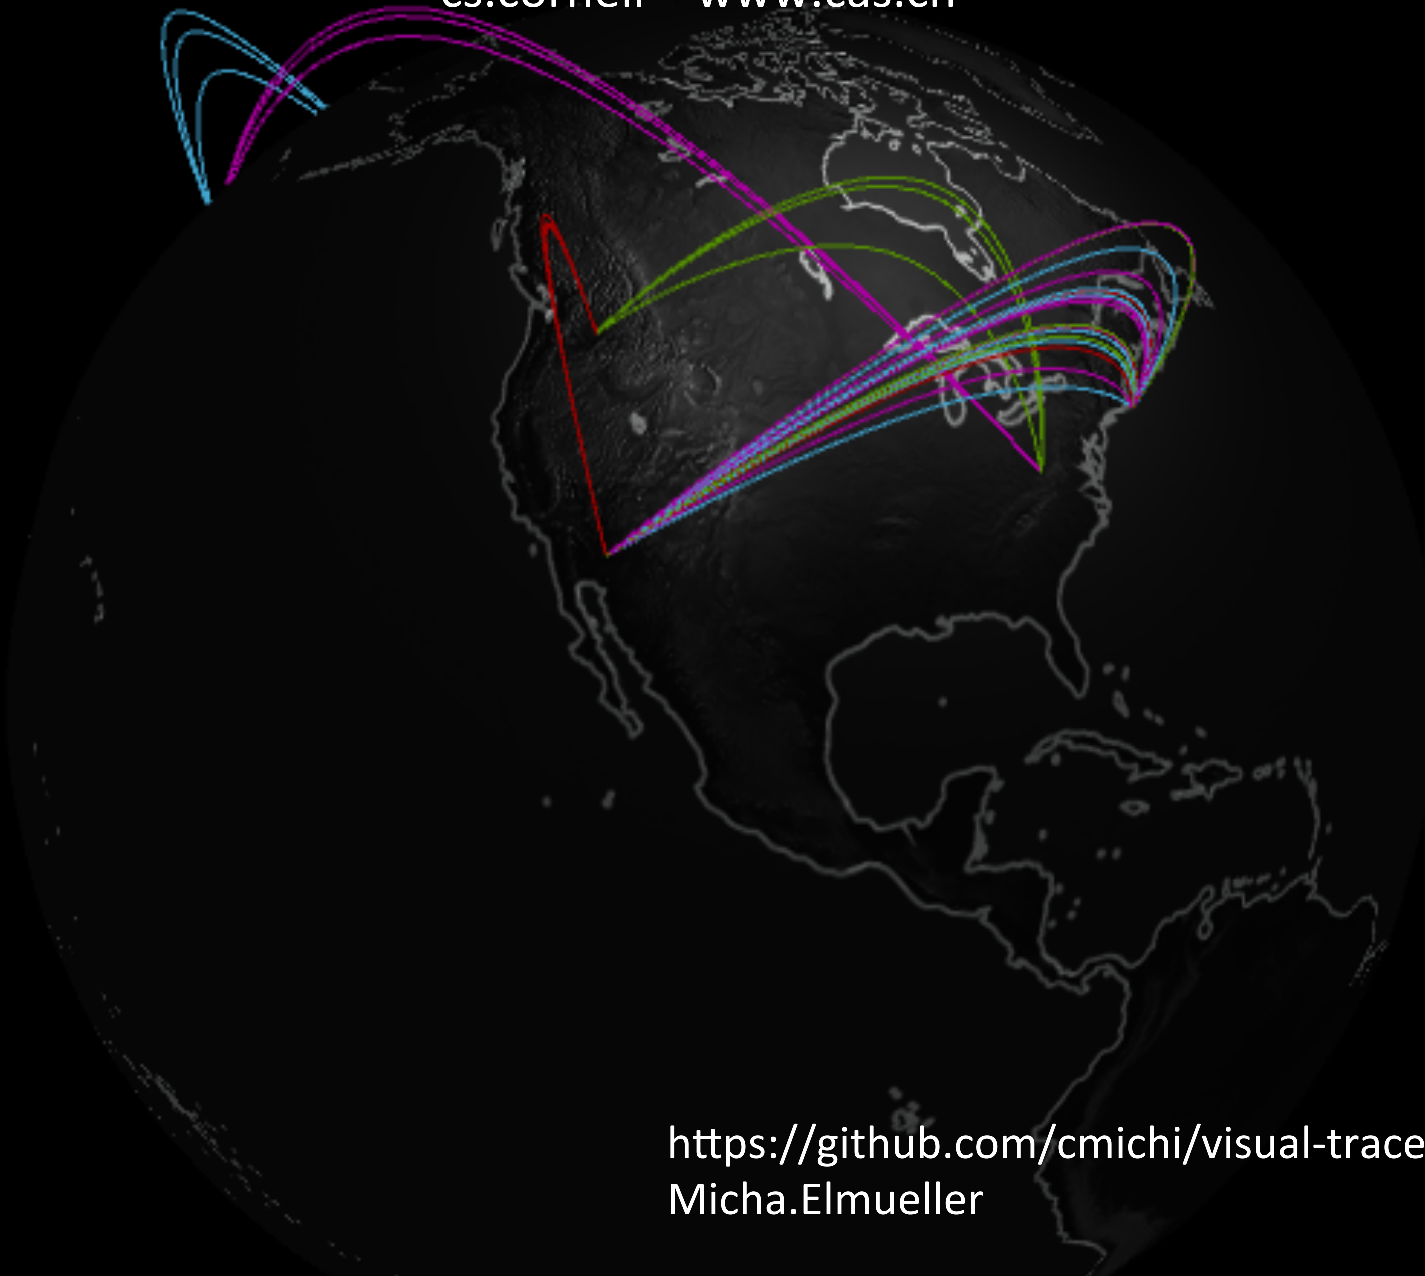

https://github.com/cmichi/visual-traceroute
Micha.Elmueller

Traceroute visualization

cs.cornell – www.cas.cn

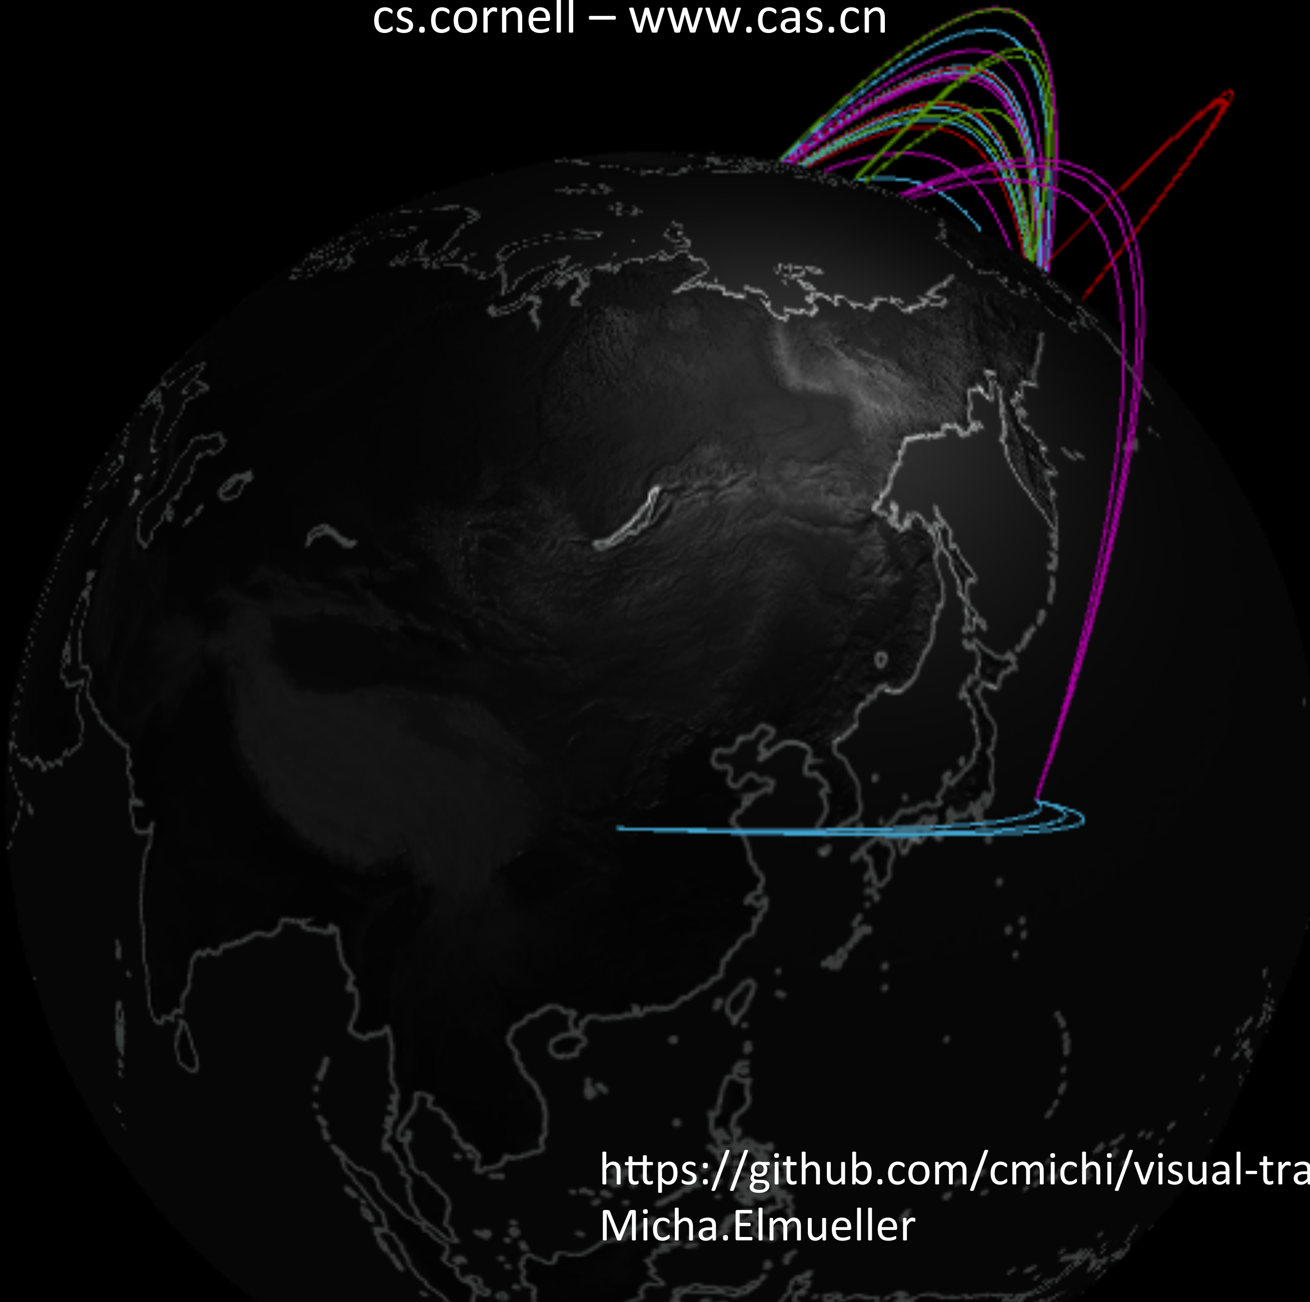

https://github.com/cmichi/visual-traceroute
Micha.Elmueller

MEASURING THE GROWTH OF THE INTERNET

The Growth of The Internet

- 4 billion IPv4 addresses depleted….

2-Layer Internet Routing Architecture

- Interdomain routing: Border Gateway Protocol (EBGP) RFC4271 (Jan 2006)
- Internal routing (Interior Gateway Protocol, IBGP)
  - Domain specific routing policies
  - Domains = Autonomous Systems
    - Not every domain network has a ASN (ASNs express distinct interdomain routing policies)

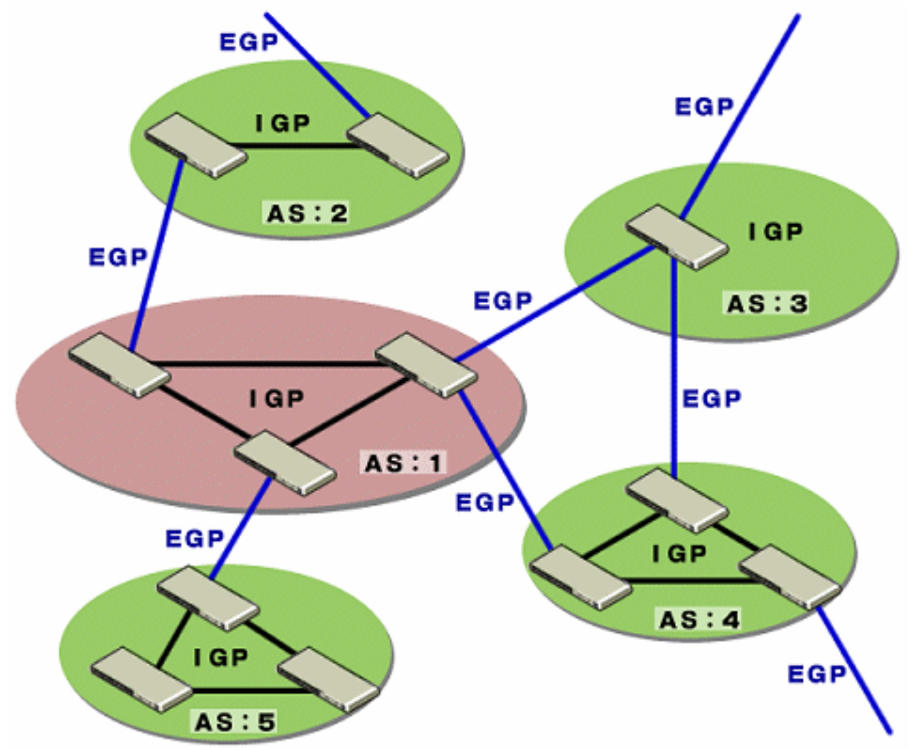

Image source: http://www.atmarkit.co.jp/fnetwork/rensai/iprt01/iprt01.html

[source: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-1autonomous_system_numbers.html, Geoff Huston, APNIC]

Autonomous Systems (AS)

"The classic definition of an Autonomous System is **a set of routers under a single technical administration, using an interior gateway protocol (IGP) and common metrics to determine how to route packets within the AS, and using an inter-AS routing protocol to determine how to route packets to other ASes.** Since this classic definition was developed, it has become common for a single AS to use several IGPs and, sometimes, several sets of metrics within an AS." [source: RFC 4271 (BGP-4, January 2006) Autonomous System (AS)]

AS Number:

Historical: 16 bit number (65,536 values) [expected to be exhausted by 5 Jul 2014, see http://www.potaroo.net/tools/asns/]

Since January 2007: 32 bit number (4,294,967,296 values)

[source: http://www.cisco.com/web/about/ac123/ac147/archived_issues/ipj_9-1/autonomous_system_numbers.html, Geoff Huston, APNIC]

Internet Registries

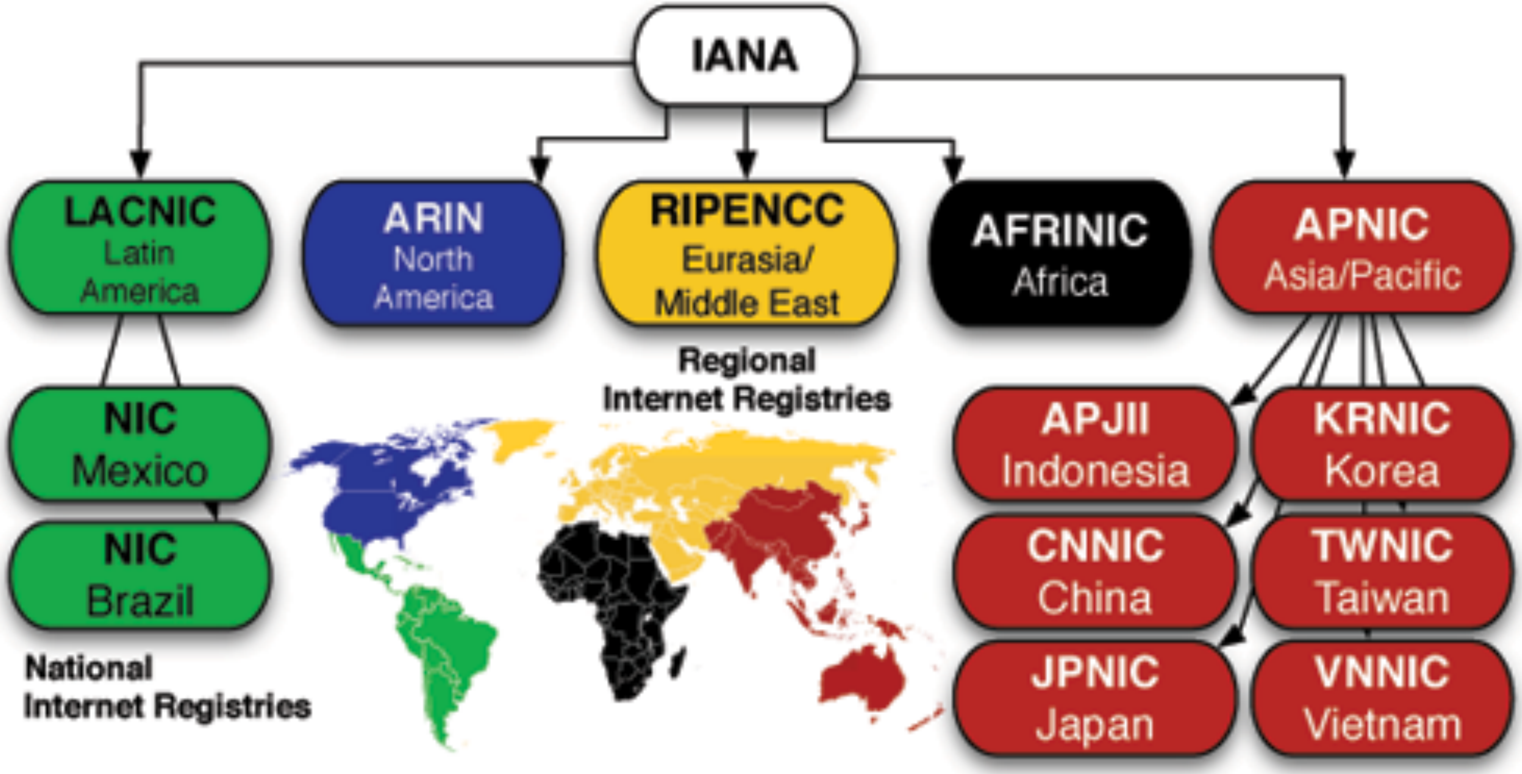

16 bit AS numbers (status as of 2012)

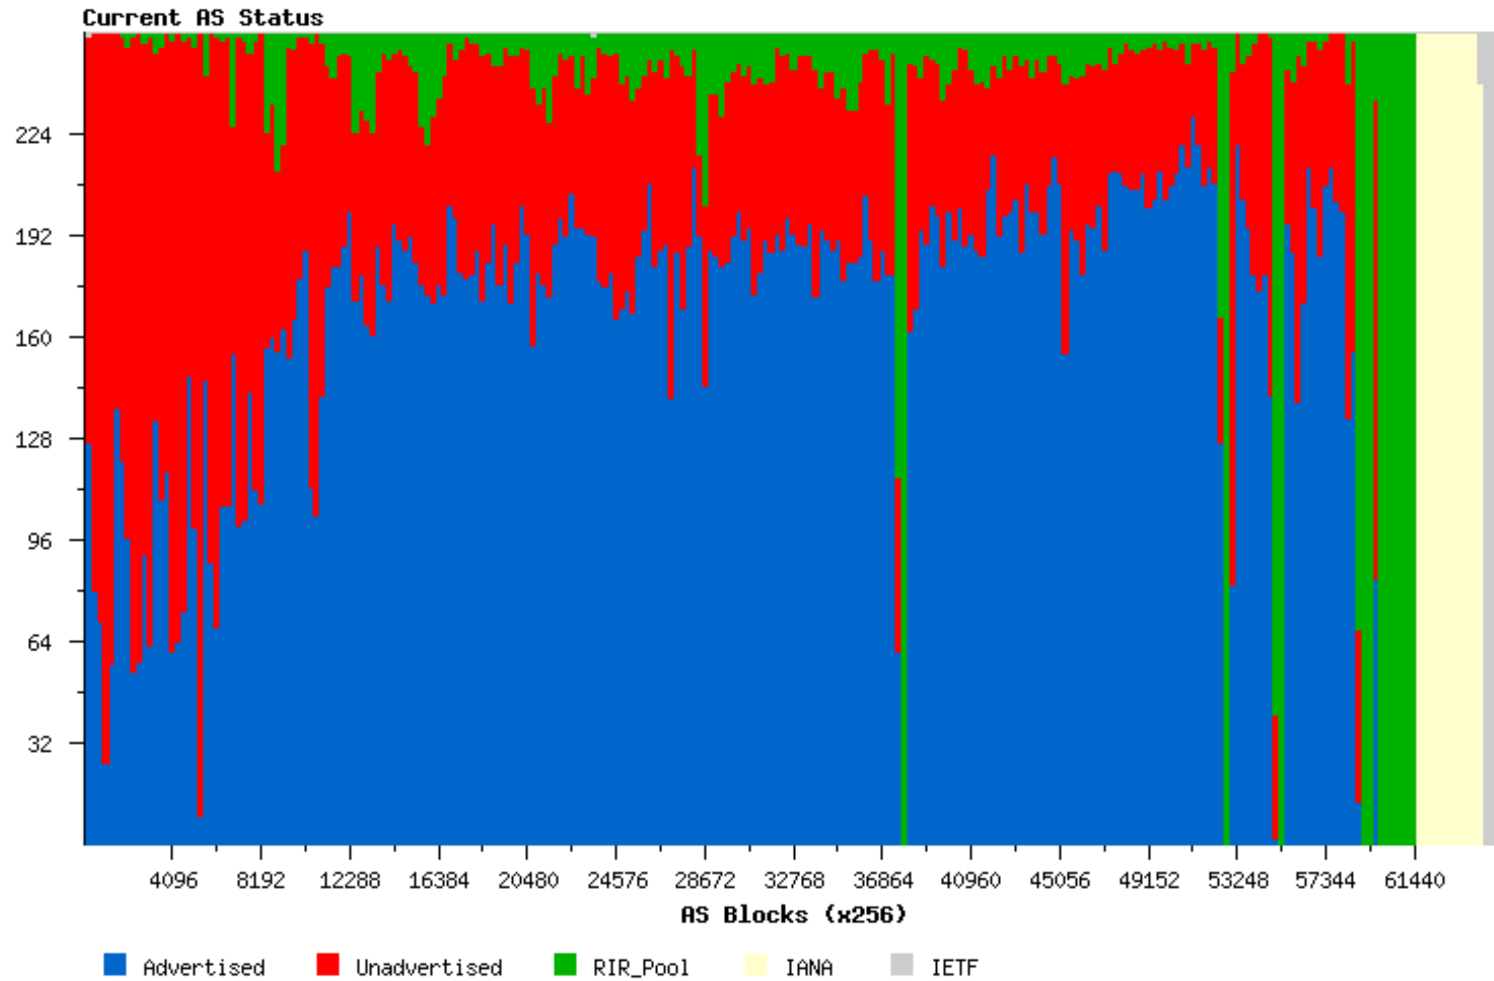

Figure 5 - AS Number Status

29-Aug-2012 07:50 UTC

16 bit AS numbers (evolution)

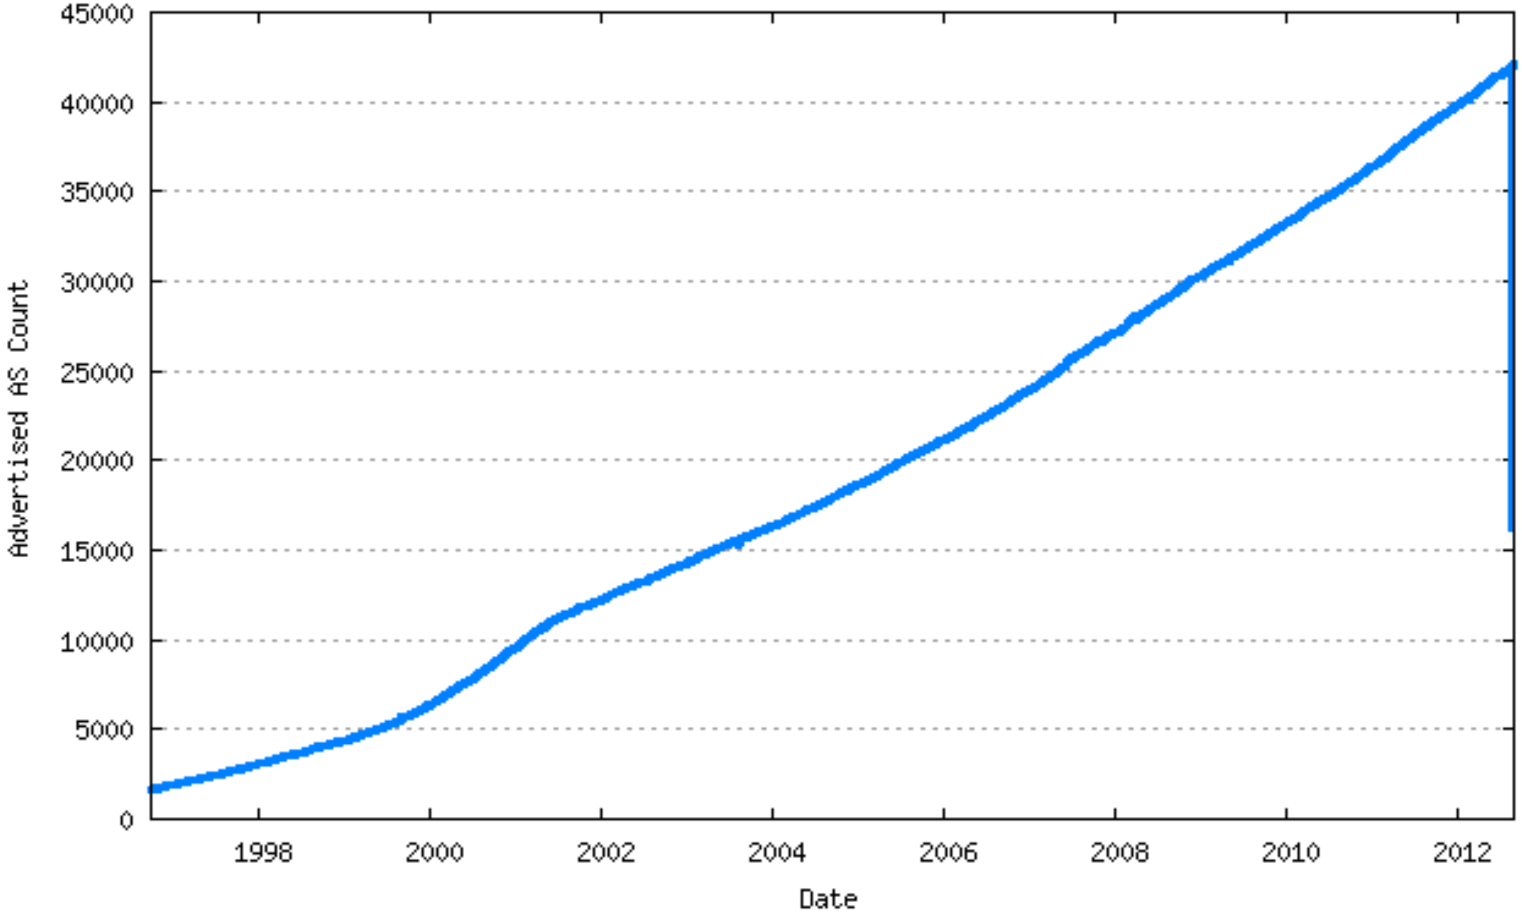

Figure 11 - Advertised AS Count

29-Aug-2012 07:50 UTC

Border Gateway Protocol Routing Tables

- Essential to the decentralized nature of the Internet
- Uses Transfer Control Protocol (TCP), port 179
- Manually configured connections to peer (neighbors)
- Peers exchange routing information
- BGP Routing tables:
  - Paths (AS numbers) to destinations (only one per destination in forwarding table)
  - Broadcast ‘update’ to neighbors, add own AS number as prefix
  - Updates from neighbors stored in Routing Information Base (RIB)
  - Routers import no loops

Routing Table Growth 1990-2014

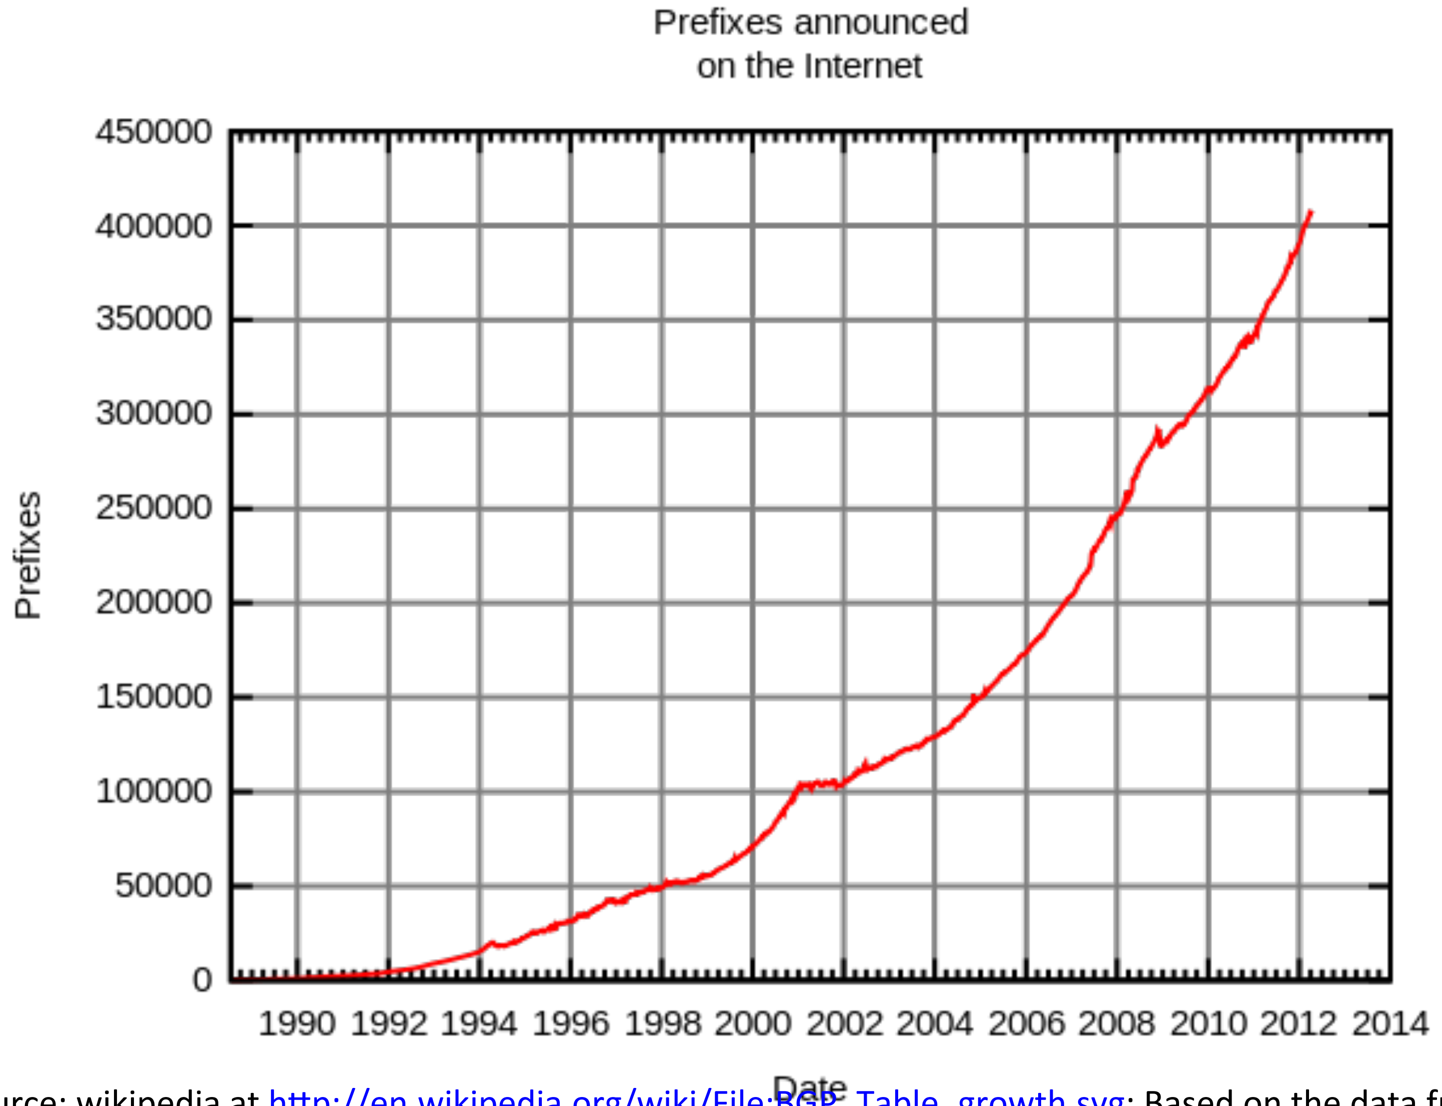

Source: wikipedia at http://en.wikipedia.org/wiki/File:BGP_Table_growth.svg; Based on the data from the CIDR report by T. Bates and al. Seee also http://bgp.potaroo.net/

AS Network Visualization

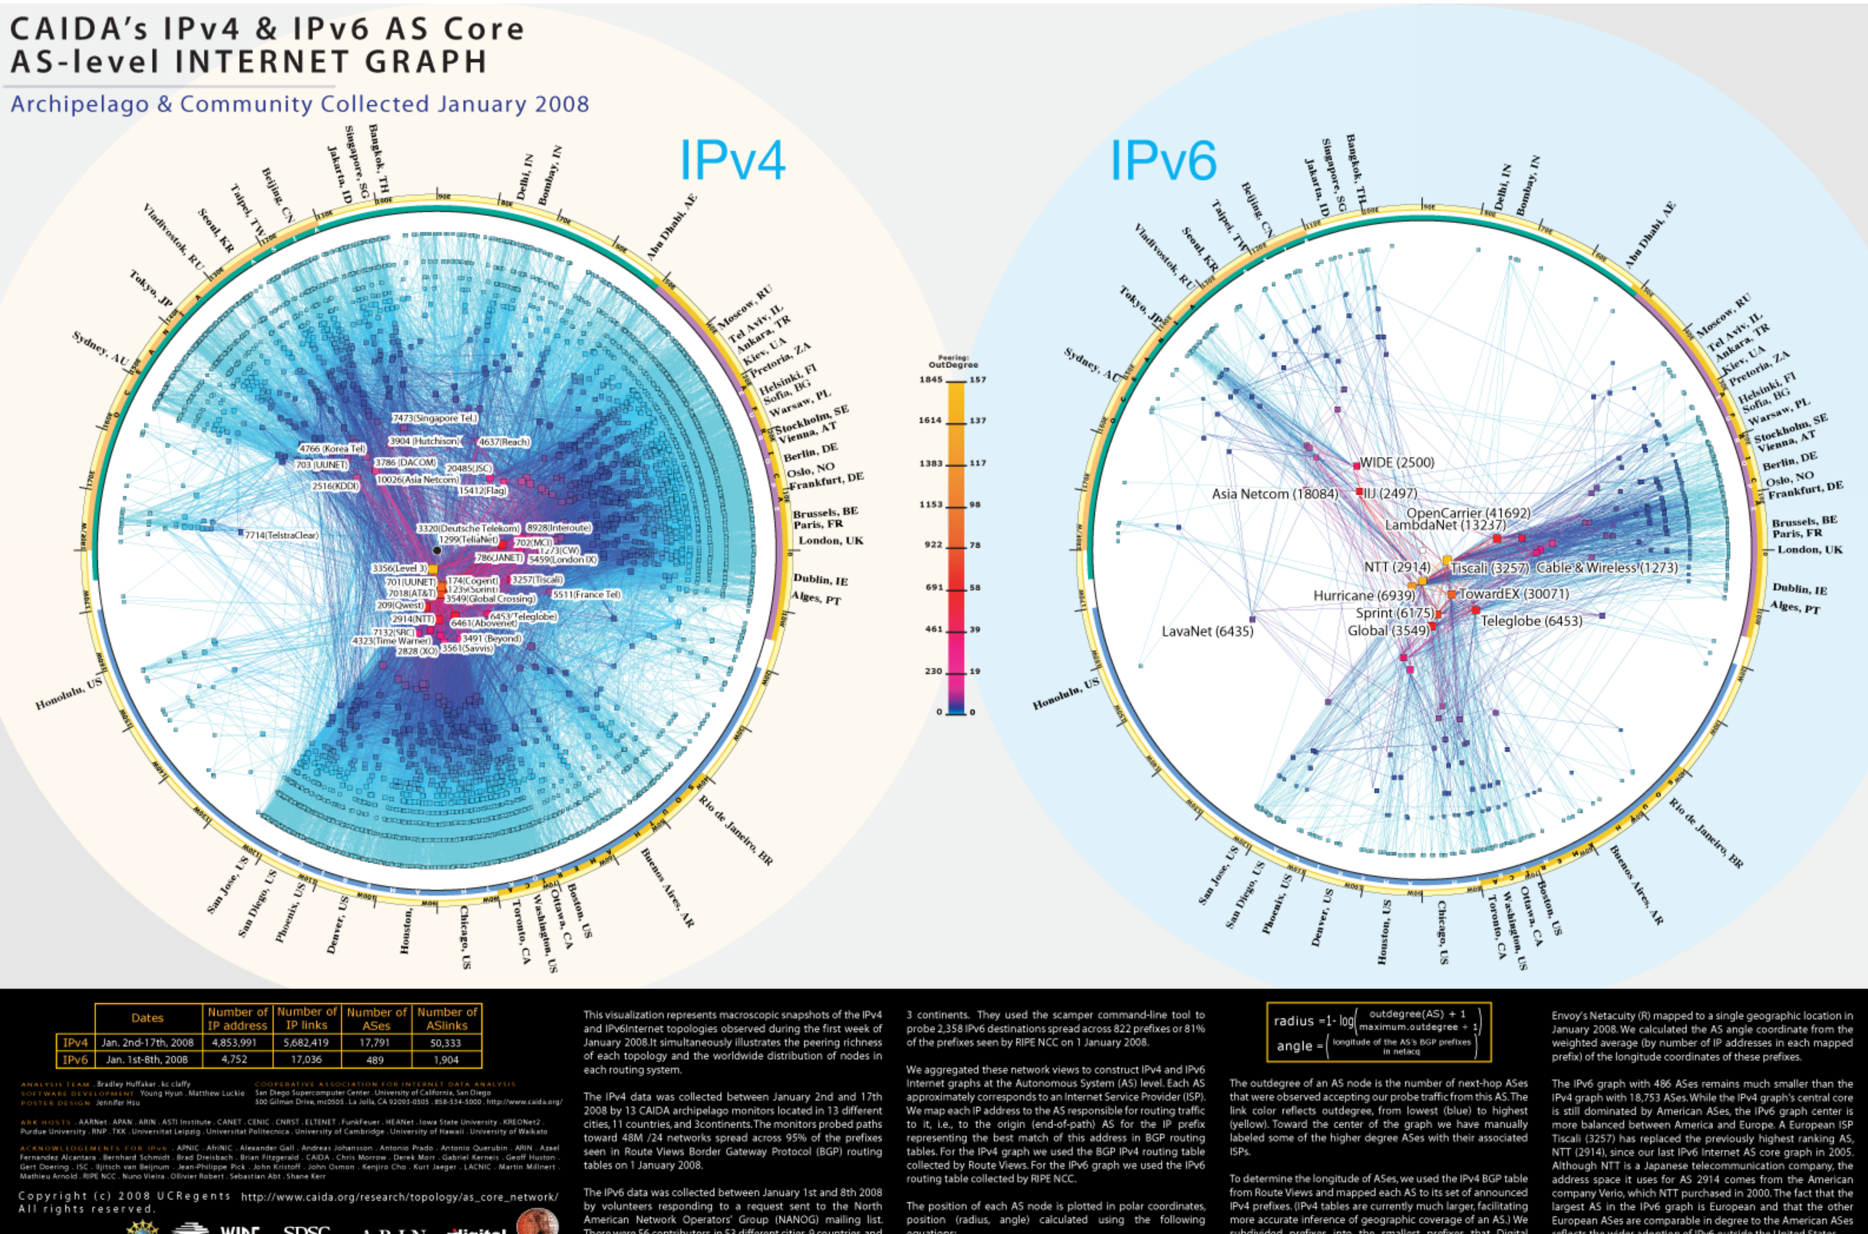

[source Cooperative Association for Internet Data Analysis http://www.caida.org/]

Historical AS Network Evolution

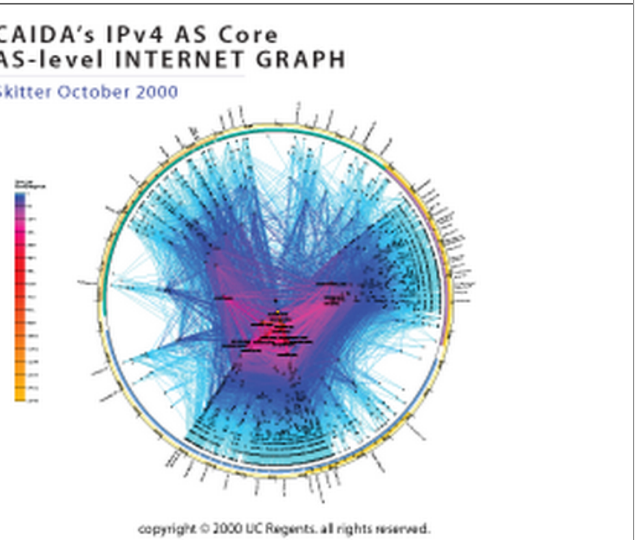

26,773 IP addresses, 1,007,723 IP links, 52% of obally routable network prefixes; 7,563 ASes 31% of A Ses present in RouteViews BGP tables n October 15, 2000), 25,005 peering sessions.
ata source: skitter

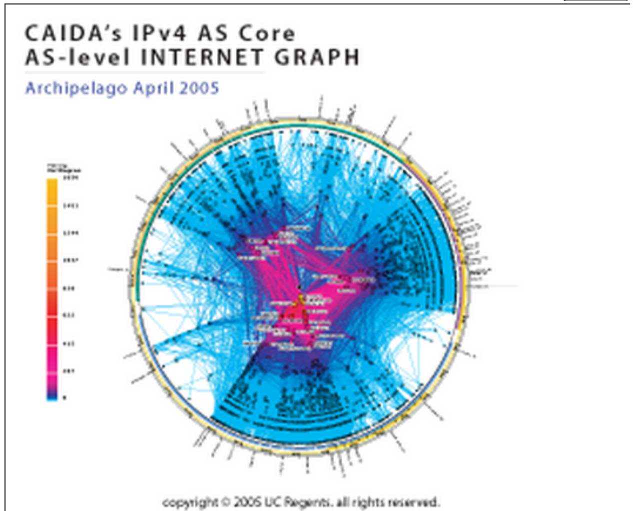

926,201 IP addresses, 2,000,796 IP links, 865,000 destinations, 50% of globally routable network prefixes; 10,131 ASes (78% of ASes present in RouteViews BGP tables on Apr 17, 2005), 29,191 peering sessions.
Data source: Ark

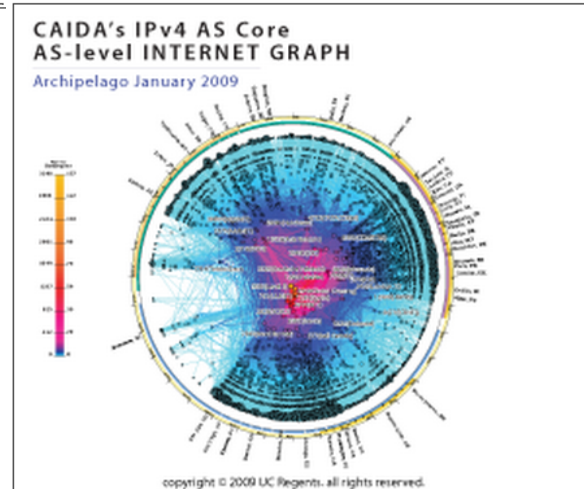

4,853,911 IP addresses, 5,682,419 IP links, 7.4 million destinations, 95% of globaly routable network prefixes, 17,791 Autonomous Systems (ASes), 50,333 peering sessions. Data collected from 33 monitors located in 30 countries on 5 continents.
Data source: Ark

[source Cooperative Association for Internet Data Analysis http://www.caida.org/]

Issues with measuring Internet Topology

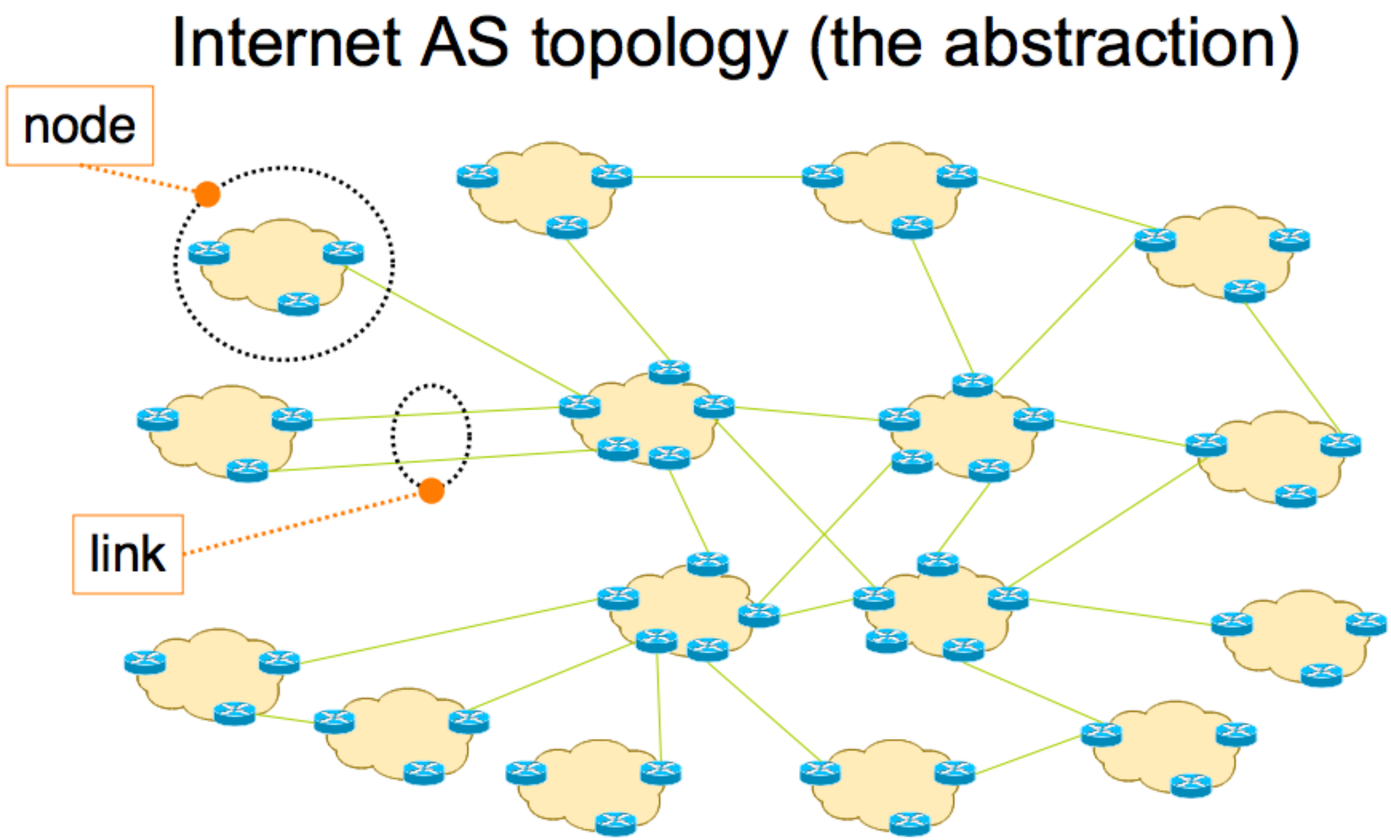

Source: Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, Lixia Zhang, "In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure", ACM SIGMETRICS, Annapolis, Maryland, June 2008.

Issues with measuring Internet Topology

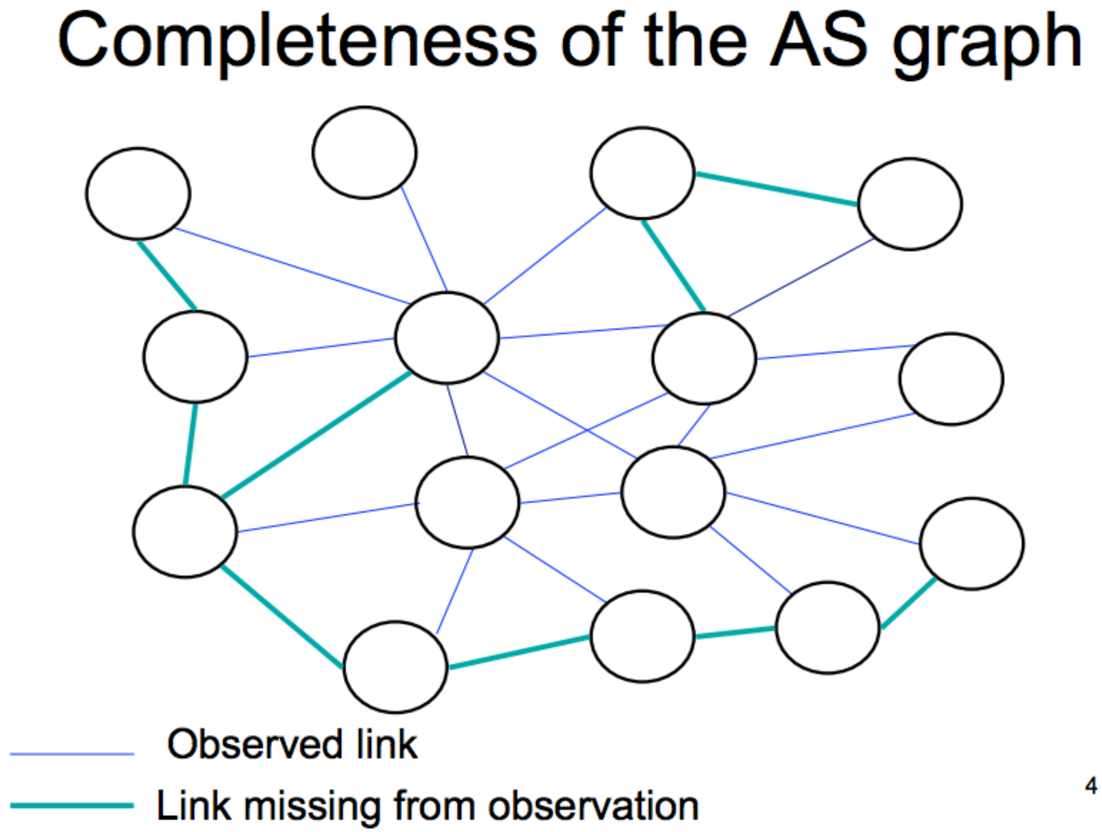

Source: Ricardo Oliveira, Dan Pei, Walter Willinger, Beichuan Zhang, Lixia Zhang, "In Search of the Elusive Ground Truth: The Internet's AS-level Connectivity Structure", ACM SIGMETRICS, Annapolis, Maryland, June 2008.

Words of caution about measuring Internet topology via BGP tables or using traceroute

M. Roughan, W. Willinger, O. Maennel, D. Perouli, and R. Bush, 10 Lessons from 10 Years of Measuring and Modeling the Internet's Autonomous Systems, IEEE Journal on Selected Areas in Communications, Vol. 29, No. 9, pp. 1-12, Oct. 2011.

Ricardo Oliveira, Dan Pei , Walter Willinger, Beichuan Zhang, Lixia Zhang, "The (in)Completeness of the Observed Internet AS-level Structure", IEEE/ACM Transactions on Networking, February 2010.

Types of Internet Traffic

- http://www.research.att.com/techdocs/TD_100193.pdf

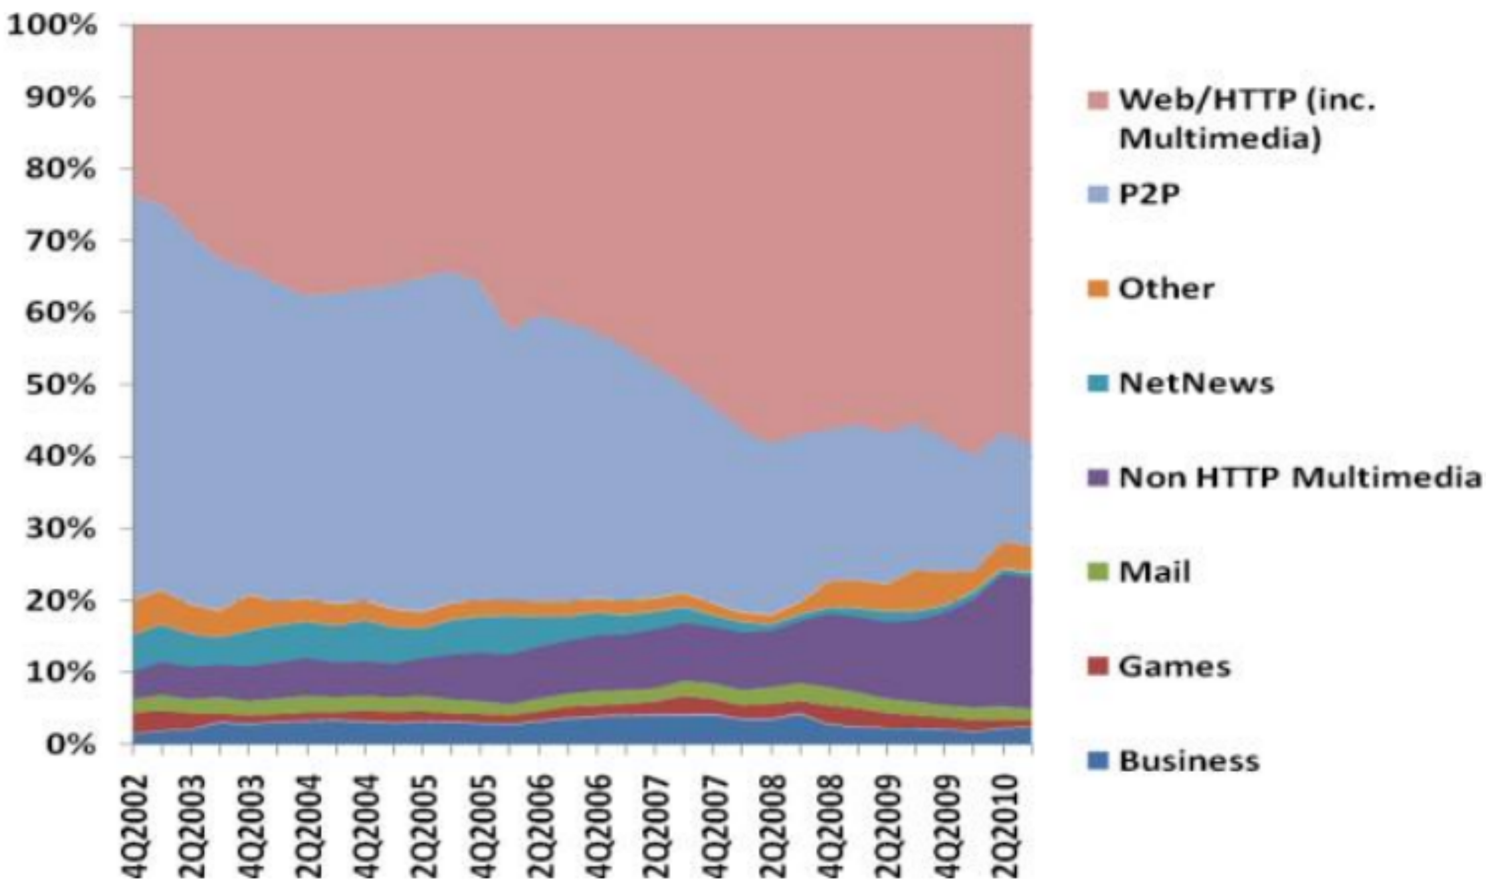

Types of Internet Traffic

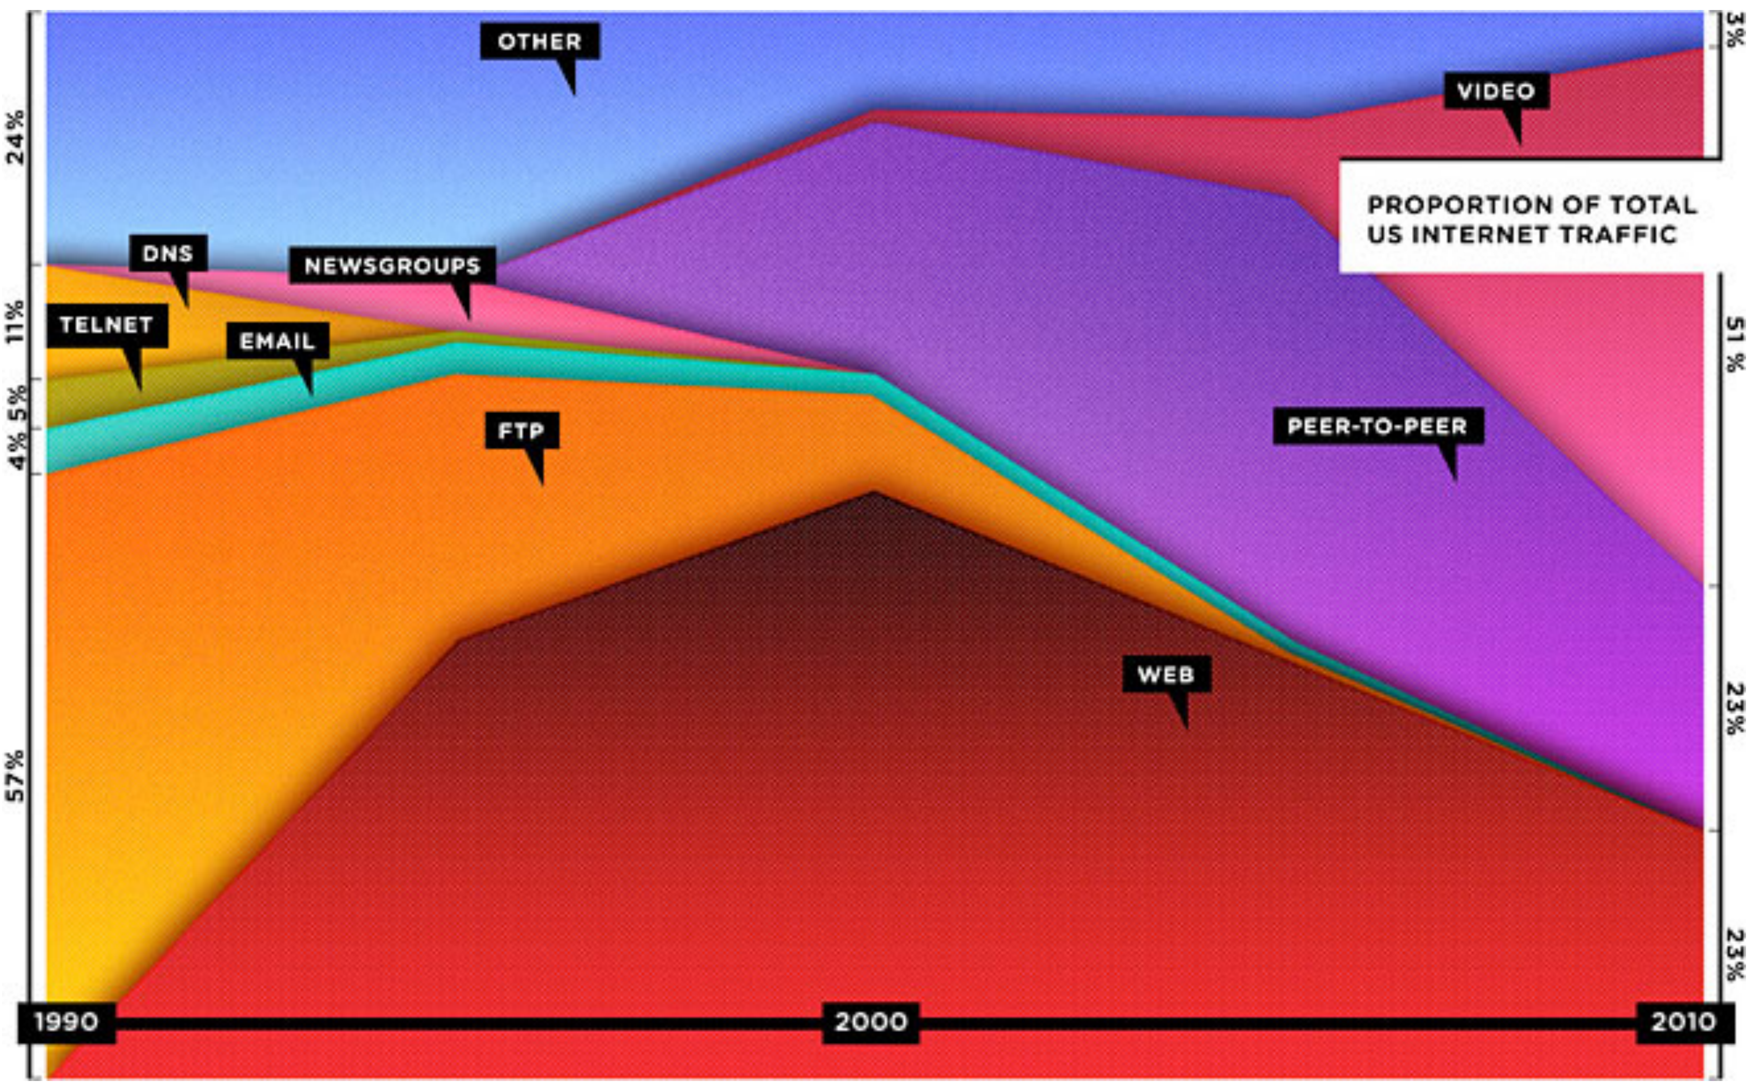

[Source: Wired Magazin 2011]

OUTLOOK:
WEEK 6 (CROSS CUTTING ISSUE)

Net Neutrality & Openness

Tim Berners Lee

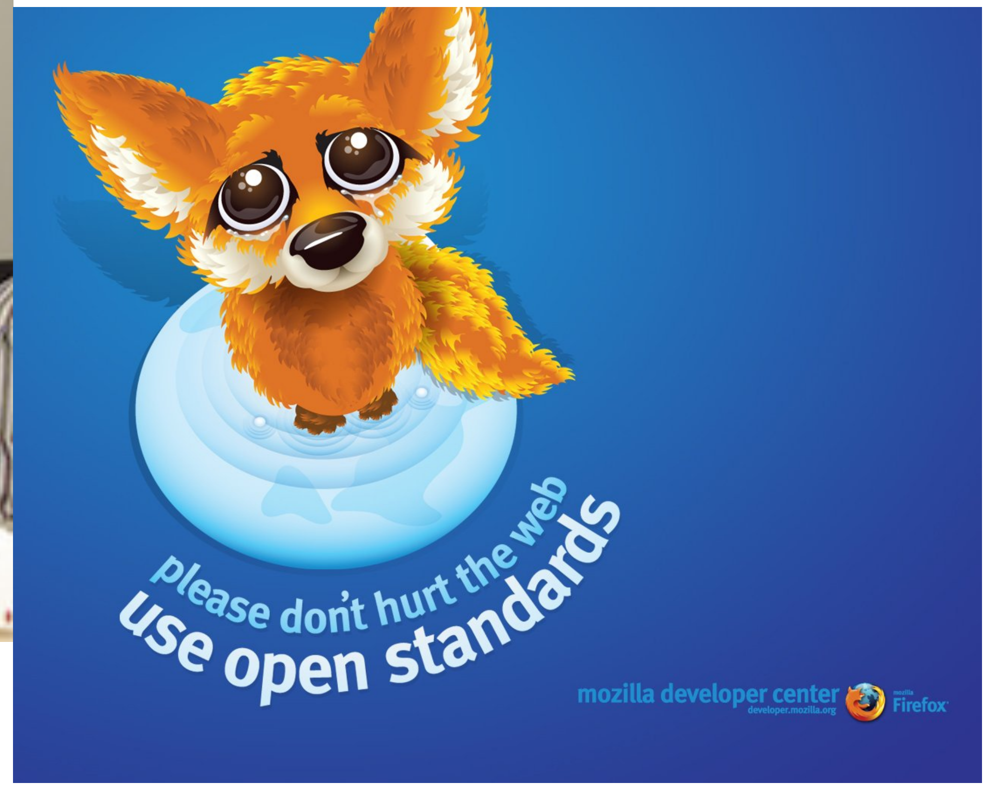

Internet Censorship

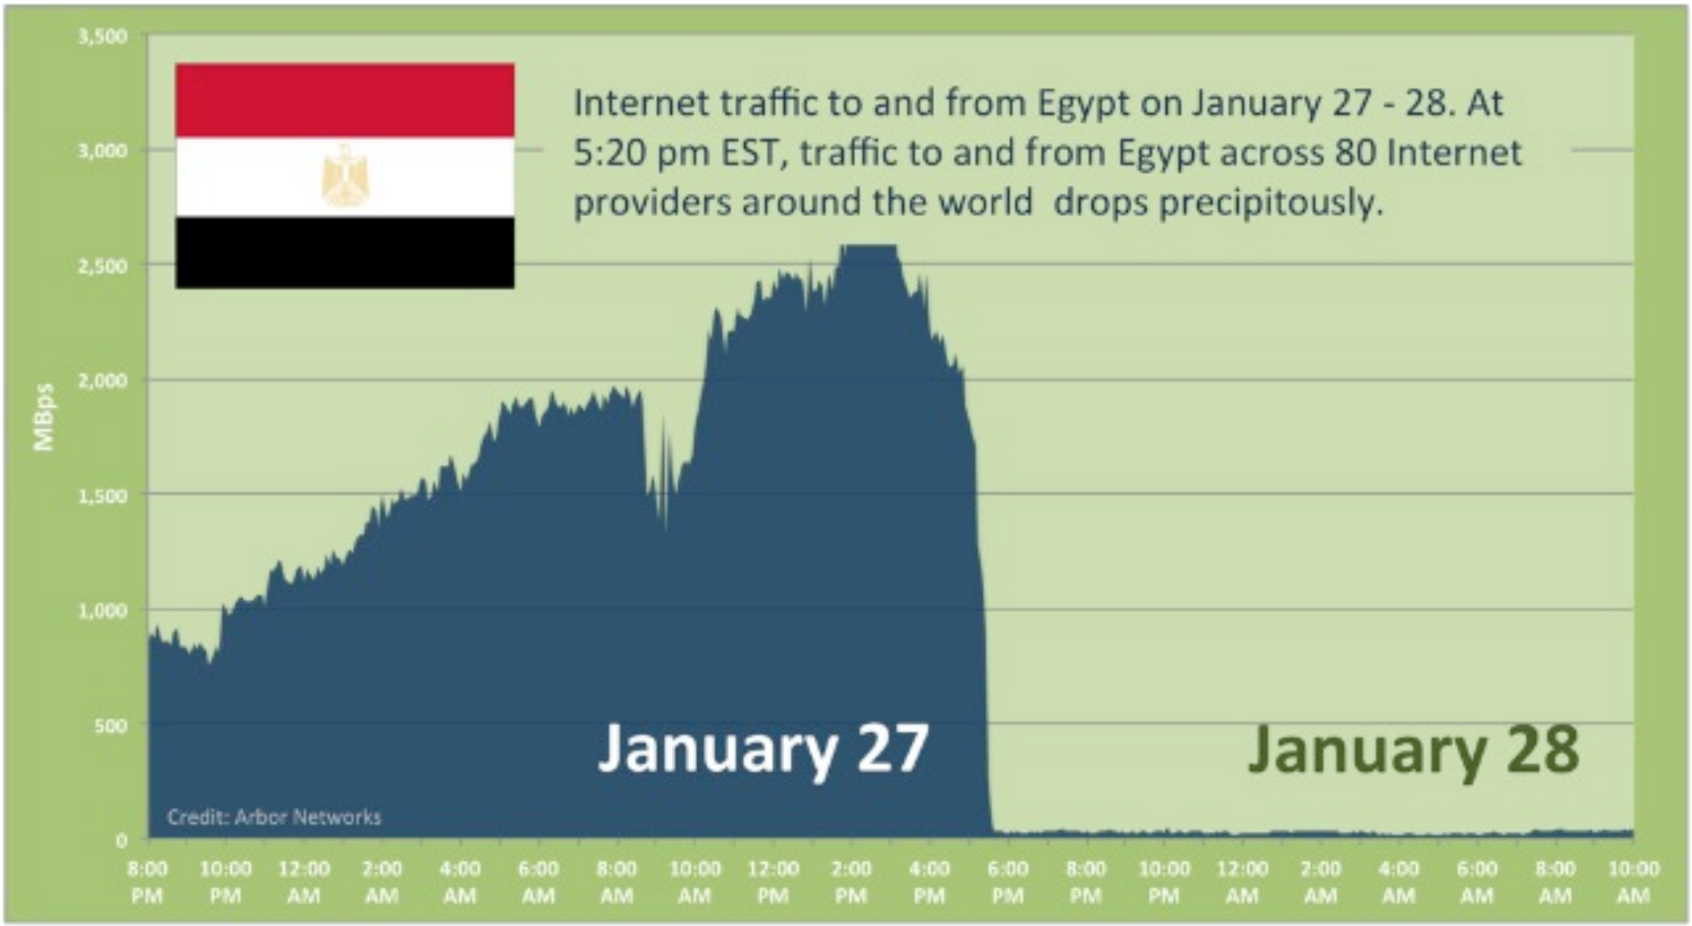

Source: arbor networks http://ddos.arbornetworks.com/2011/01/egypt-loses-the-internet/

Internet Surveillance

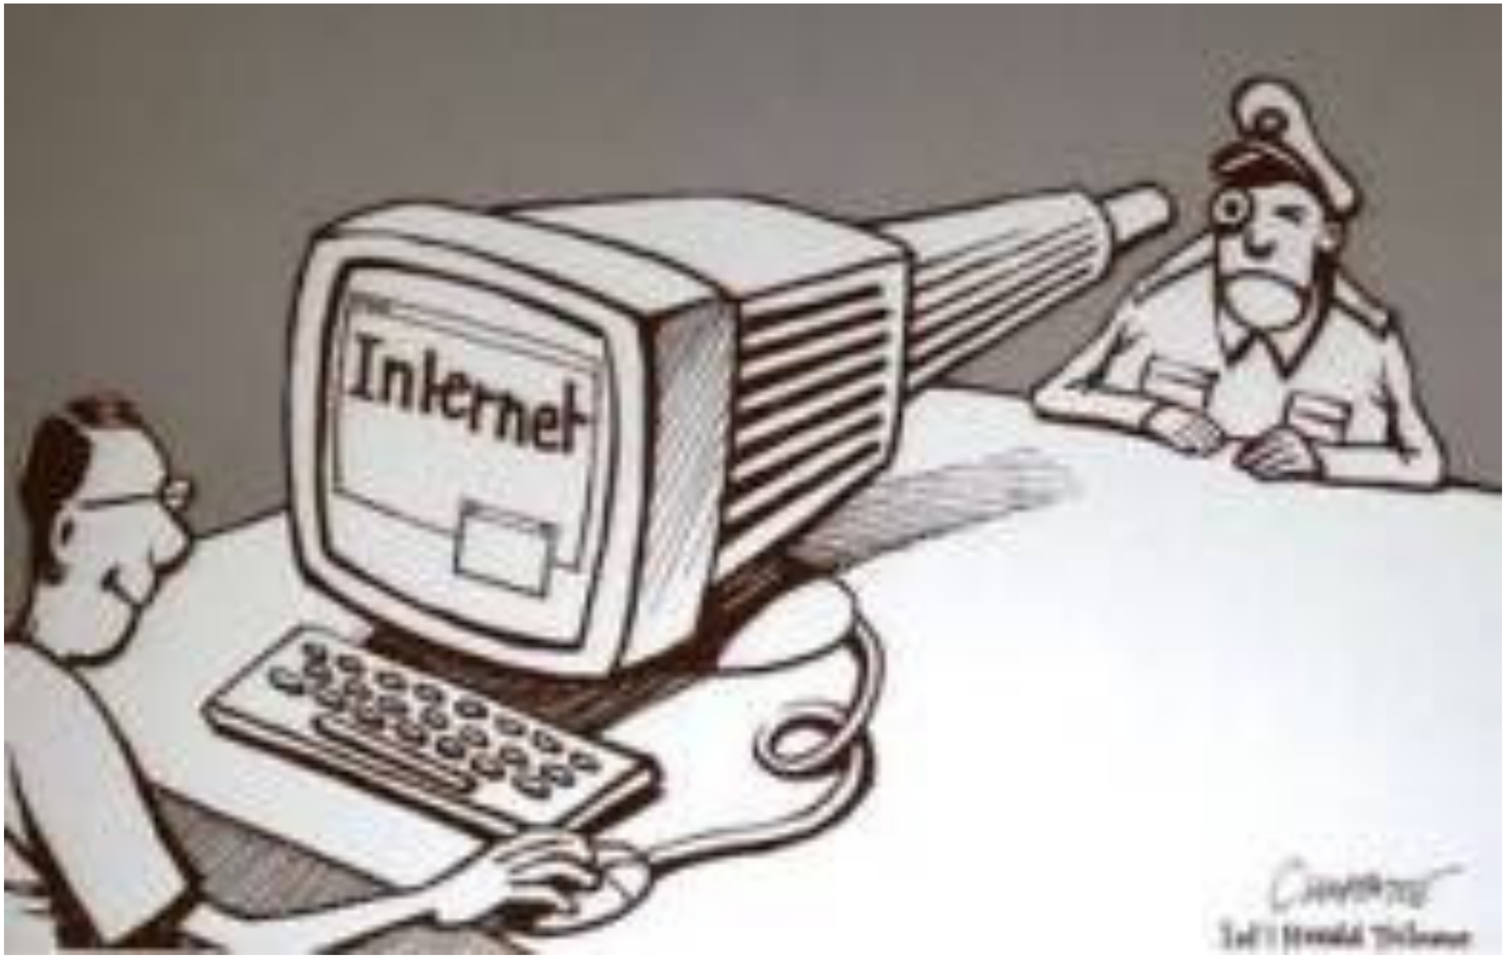

Cartoon by: Patrick Chappatte (International Herald Tribune)

Cast your vote on piazza:

Cross Cutting Issue for Week 6:

1. Net Neutrality
2. Internet Censorship
3. Internet Surveillance

Next Week:

- Week 3: History and Architecture of The World Wide Web
- Reminders:
  - Online course communication exclusively through piazza
  - Homework 1 due on Sunday via CMS (check whether you are on CMS otherwise email me on tav6)
  - Course website at http://www.infosci.cornell.edu/Courses/info4302/2012fa/