Short Introduction to IPv6

Martin Schütte

```
Network Working Group                                         S. Deering
Request for Comments: 2460                                         Cisco
Obsoletes: 1883                                                R. Hinden
Category: Standards Track                                          Nokia
                                                           December 1998


                Internet Protocol, Version 6 (IPv6)
                           Specification

Status of this Memo

   ...

Abstract

   This document specifies version 6 of the Internet Protocol (IPv6),
   also sometimes referred to as IP Next Generation or IPng.
```

The Hourglass Internet Architecture

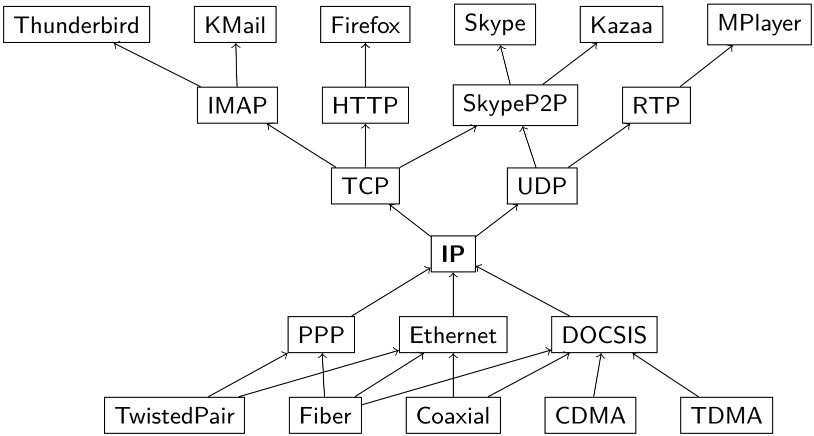

by Constantine Dovrolis

Addresses

IPv4: 32 bits

11001011000000000111000100010111

11001011	00000000	01110001	00010111
0xcb	0x00	0x71	0x17
203	0	113	23

203.0.113.23

Addresses

IPv6: 128 bits

```
0010000000000001000001100011100000001000000001110000000000111010
0000001000000101010111011111111111111110111110110011100000010010
```

```
0010000000000001  0000011000111000  0000100000000111  0000000000111010
0000001000000101  0101110111111111  1111111011111011  0011100000010010
```

```
0x2001   0x0638   0x0807   0x003a
0x0205   0x5dff   0xfefb   0x3812
```

`2001:0638:0807:003a:0205:5dff:fefb:3812`

Notation

```
2a01:0198:0200:01e4:0000:0000:0000:0002
= 2a01:198:200:1e4:0:0:0:2
= 2a01:198:200:1e4::2
```

```
2001:0db8:0000:0000:0001:0000:0000:0001
= 2001:db8:0:0:1:0:0:1
= 2001:db8::1:0:0:1
```

Ports: `[2001:db8::1:0:0:1]:443`
`http://[2001:db8::1:0:0:1]:443/index.html`

Address Structure

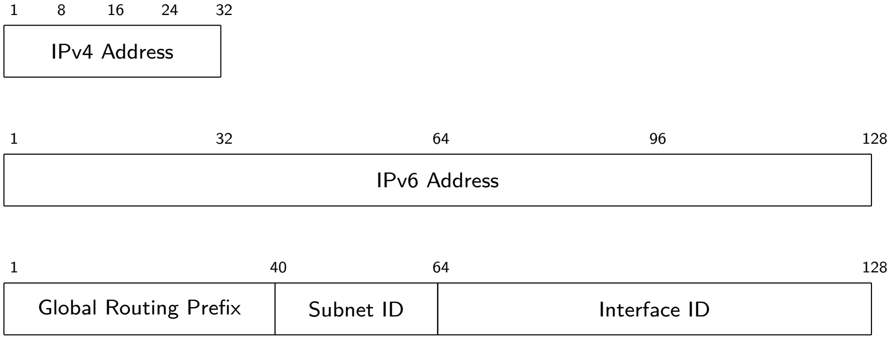

Address Allocation

1	23	32	48	64	128
IANA	RIPE	LIR	Site	Interface ID	

Example: University Server `2001:638:807:3a:205:5dff:fefb:3812`

1	23	32	48	64	128
IANA `2001:06`	RIPE `38:`	DFN `0807:`	Uni `003a:`	Interface ID `0205:5dff:fefb:3812`	

Address Usage

Example: Comcast (USA)

- 20 M customers
- avg. 2.5 set-top boxes (STB) per customer
- 2 IPs per STB

$\Rightarrow$ $20\text{M} \times 2.5 \times 2 = 100\,\text{M}$ IPs

Comparison:
10.0.0.0/8 has $2^{24} \approx 16.5$ M IPs

Address Usage

Hypothetical example: China Mobile

- 740 M customers
- avg. 2 devices per customer
- 2 IPs per device

$\Rightarrow$ $740\text{M} \times 2 \times 2 \approx 3000\,\text{M}$ IPs

Comparison:
$0/1$ has $2^{31} \approx 2150$ M IPs
IPv4 has $2^{32} \approx 4300$ M IPs

Some Special Address Ranges

`::`	unspecified
`::1`	localhost
`::ffff:192.0.2.1`	IPv4-mapped
`2000::/3`	unicast
`fe80::/10`	link-local
`fd00::/8`	site-local
`ff00::/8`	multicast
`ff02::1`	multicast, all hosts
`ff02::2`	multicast, all routers
`ff02::101`	multicast, all NTP servers
`ff05::1:3`	multicast, all DHCP servers

Scope

Multicast addresses have scope:

- `ff01::101` means all NTP servers on the same interface (i.e., the same node) as the sender.
- `ff02::101` means all NTP servers on the same link as the sender.
- `ff05::101` means all NTP servers in the same site as the sender.
- `ff0e::101` means all NTP servers in the Internet.

Example:

```
mschuett@zula:~$ ssh root@fe80::6670:2ff:fe4c:3a61
ssh: connect to host fe80::6670:2ff:fe4c:3a61 port 22: Invalid argument

mschuett@zula:~$ ssh root@fe80::6670:2ff:fe4c:3a61%eth0

BusyBox v1.19.4 (2013-04-07 13:28:47 CEST) built-in shell (ash)
Enter 'help' for a list of built-in commands.
...
```

IPv4 Header Format

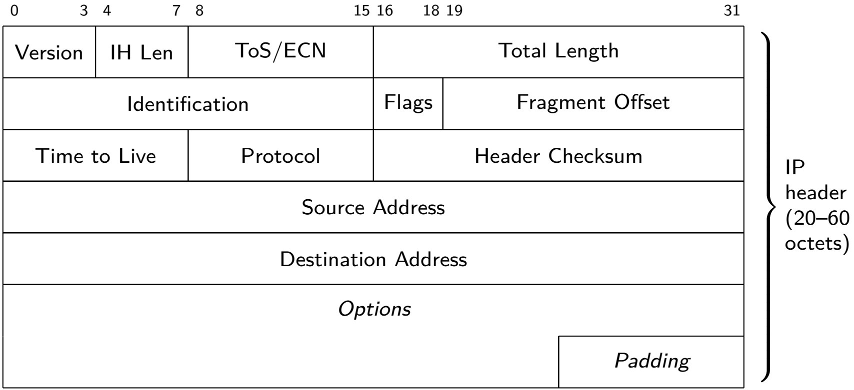

IPv6 Header Format

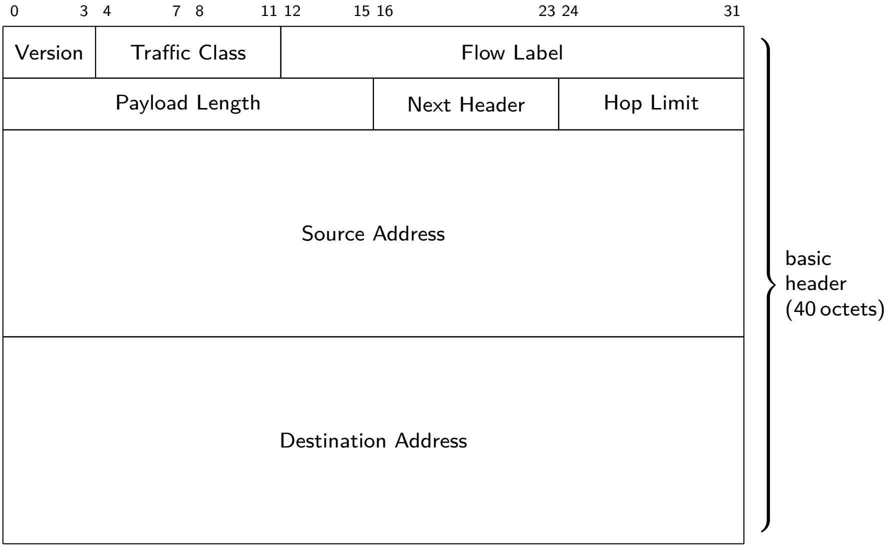

IPv6 Extension Header

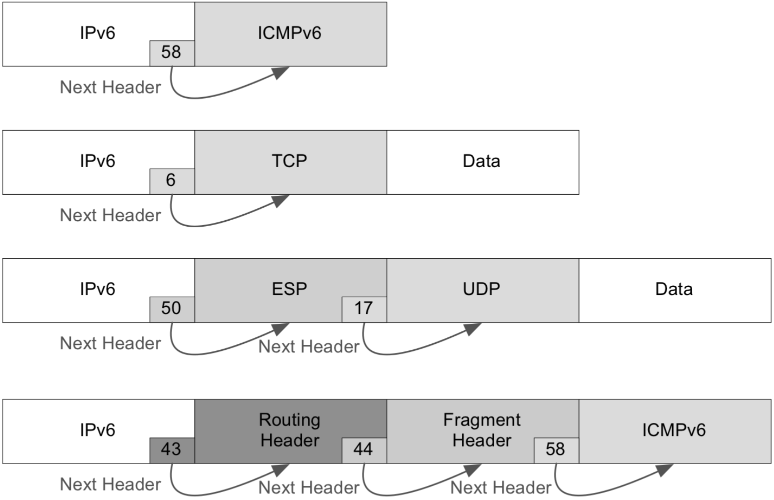

Destination/Hop-by-Hop Option Header

0 … 7	8 … 15	16 … 23	24 … 31
Next Header	Hdr Ext Len	Opt Type	Opt Len
Opt Value	. . .		

0 … 7	8 … 15	16 … 23	24 … 31
Next Header: 0x3a *ICMPv6*	Hdr Ext Len: 0x00 *8 octets*	Opt Type: 0x05 *Rtr alert*	Opt Data Len: 0x02 *2 octets*
Opt Data: 0x00 0x00 *MLD*		Opt Type: 0x01 *PadN*	Opt Data Len: 0x00 *0 octets*

ICMPv6

ICMPv6 includes:

- control messages (like ICMPv4)
- Link-Layer address resolution (like ARP)
- multicast group membership (like IGMP)
- Path MTU Discovery
- Neighbor Discovery/Autoconfiguration

Stateless Address Autoconfiguration

1. select tentative address
2. send Router Solicitation to `ff02::2`
3. join multicast groups for *all nodes* and *Solicited-Node address*
4. Duplicate Address Detection

⇒ acquired link-local IP

receive Router Advertisement(s) with network info:

- routers with lifetime and preference
- local network MTU
- global address prefix
- DHCP management flag

⇒ acquire global address

Link-Layer Address Resolution

Solicited-Node Address: `FF02:0:0:0:0:1:FF`*`XX:XXXX`*

Mapping:
Unicast `fe80::218:f3ff:fe3a:3f55`
⇒ Multicast `ff02::1:ff3a:3f55`
⇒ MAC `33:33:ff3a:3f:55`

send Neighbor Solicitation to multicast
⇒ receive Neighbor Advertisement with IP and MAC address

Note: no broadcast required(!)

Duplicate Address Detection

Solicited-Node Address: `FF02:0:0:0:0:1:FF`*`XX:XXXX`*

Mapping:
Unicast `fe80::218:f3ff:fe3a:3f55`
⇒ Multicast `ff02::1:ff3a:3f55`
⇒ MAC `33:33:ff3a:3f:55`

send Neighbor Solicitation to multicast

- receive Neighbor Advertisement with IP and MAC address ↯
- no answer, address is free ✓

Autoconfiguration and Neighbor Discovery

Design assumption: secure and trustworthy LAN

Simple Denial of Service:

1. Host Alice starts *Duplicate Address Detection*:
   "Anyone using IP *X*?"
2. Host Eve answers "I have IP *X*."
3. `goto 1`

Routing/Man in the Middle:

1. Host Eve sends ICMPv6 Redirect:
   "This is router Bob, for *google.com* please use router Eve."

DNS

IPv4: A Record — IPv6: AAAA Record

```
[mschuett@cayce] ~> dig +short quux.de a
88.198.11.141
[mschuett@cayce] ~> dig +short quux.de aaaa
2001:6f8:1138::1
[mschuett@cayce] ~> dig +short -x 2001:6f8:1138::1
mail.quux.de.
[mschuett@cayce] ~> dig -x 2001:6f8:1138::1
...
;; QUESTION SECTION:
;1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.
      8.3.1.1.8.f.6.0.1.0.0.2.ip6.arpa. IN PTR
...
```

Happy Eyeballs

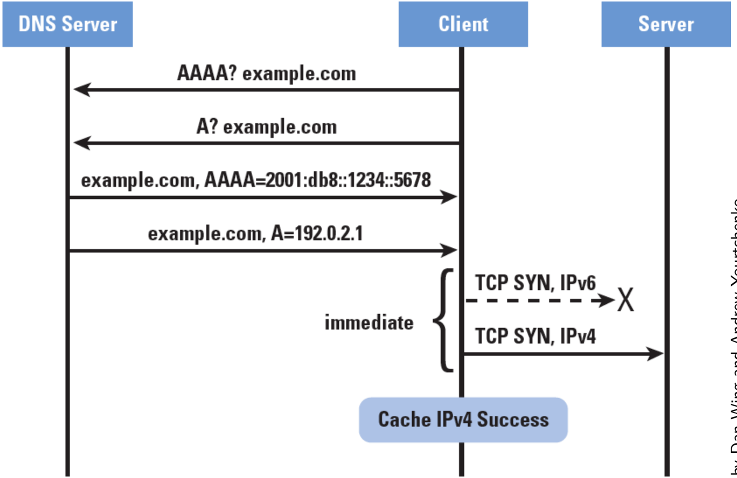

by Dan Wing and Andrew Yourtchenko

Known for Some Time

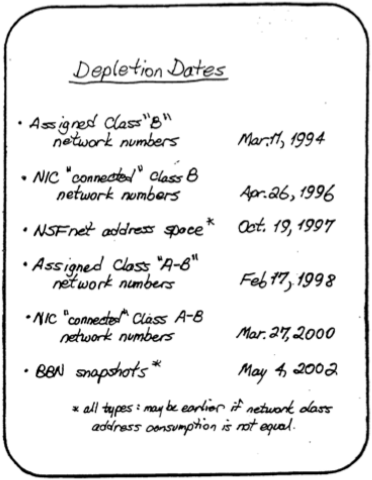

Frank Solensky's Report on Address Depletion, Proceedings of IETF 18, p. 61, Vancouver, August 1990

The Plan

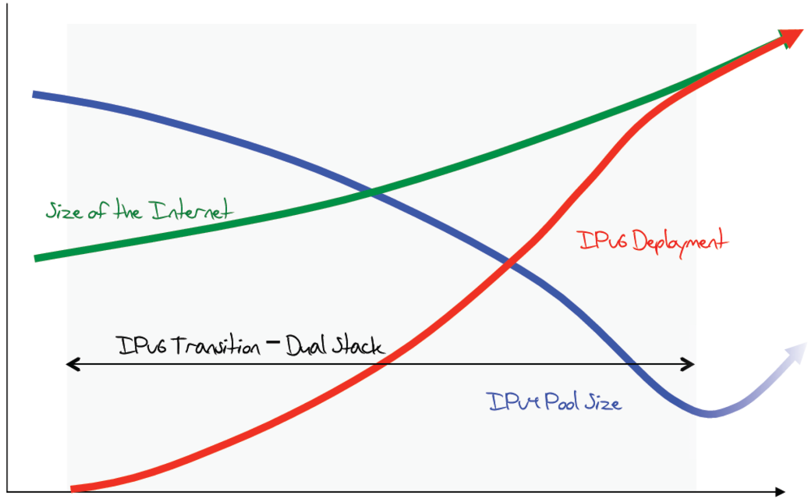

by Geoff Huston

Plan B

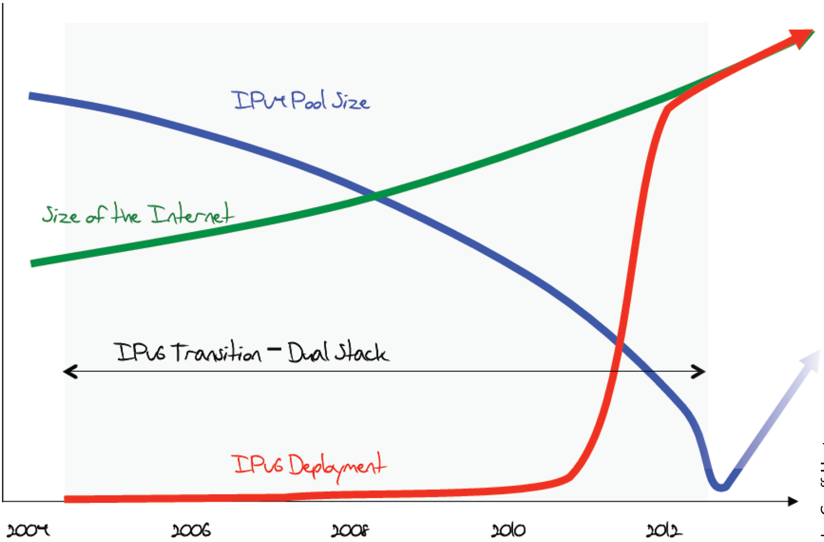

by Geoff Huston

Status Quo

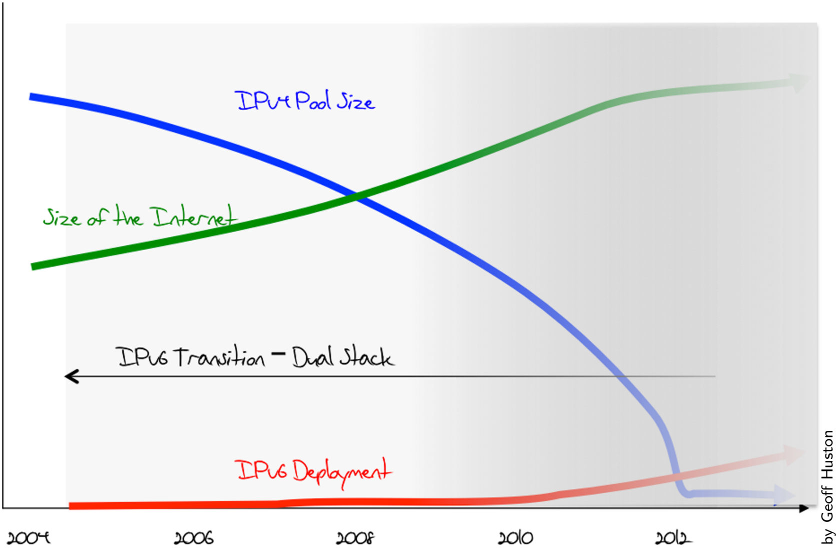
IPv4 Pool Size
Size of the Internet
IPv6 Transition – Dual Stack
IPv6 Deployment
2004
2006
2008
2010
2012
by Geoff Huston

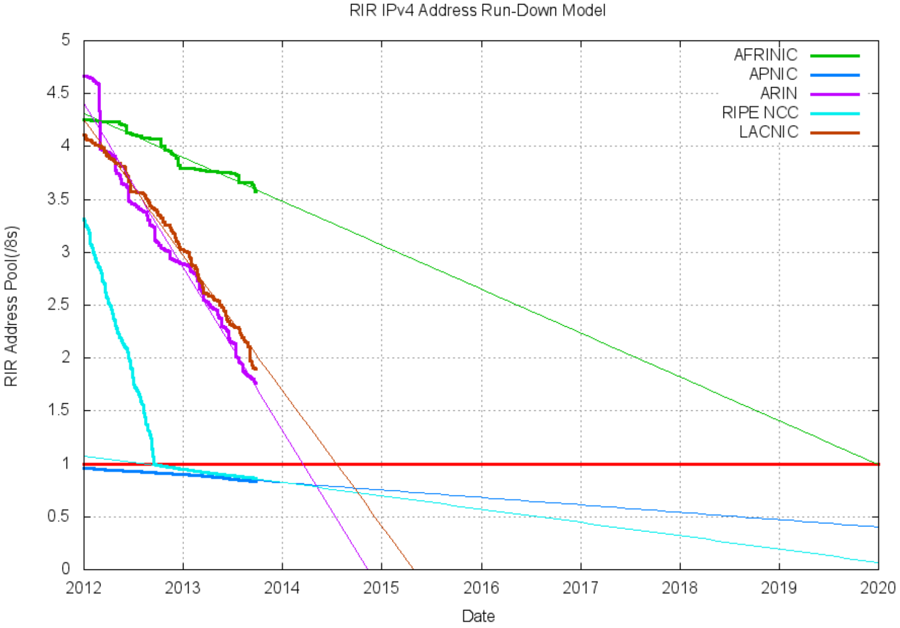
RIR IPv4 Address Run-Down Model
RIR Address Pool(/8s)
Date
AFRINIC
APNIC
ARIN
RIPE NCC
LACNIC
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
5
2012
2013
2014
2015
2016
2017
2018
2019
2020

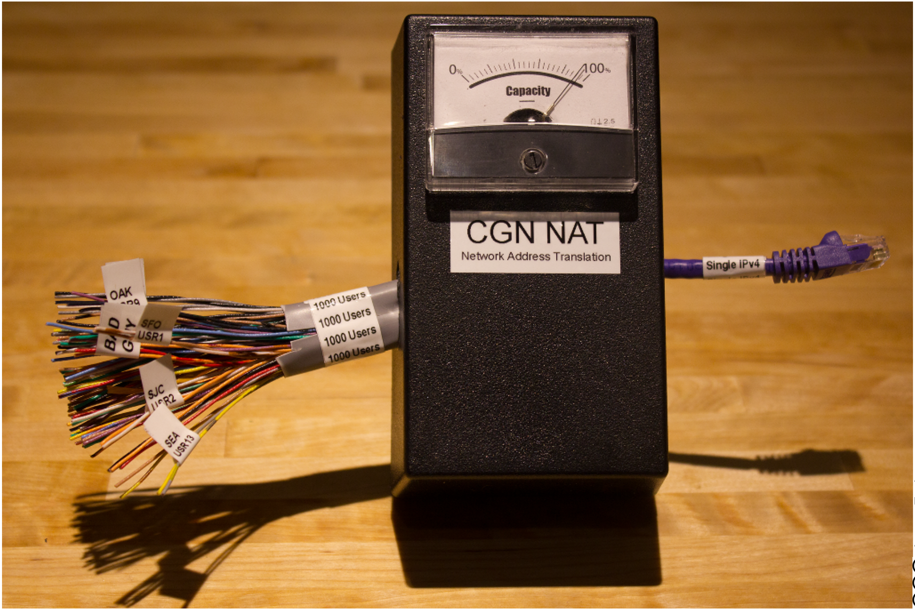
0%
100%
Capacity
CGN NAT
Network Address Translation
Single IPv4
1000 Users
1000 Users
1000 Users
1000 Users
OAK
SFO
USR1
SJC
SEA
USR13

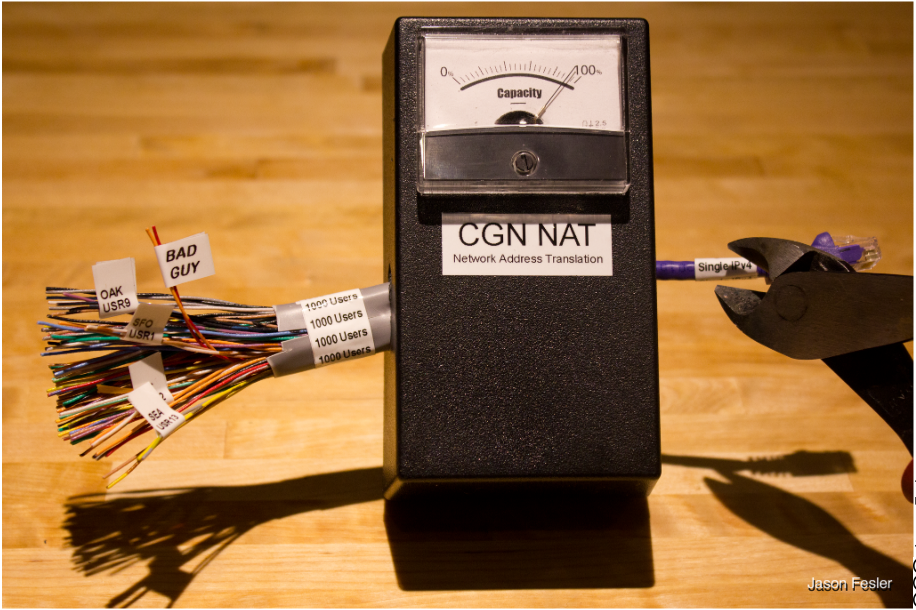
0%
100%
Capacity
CGN NAT
Network Address Translation
1000 Users
1000 Users
1000 Users
1000 Users
BAD GUY
OAK USR9
SFO USR1
SEA USR13
Single IPv4
Jason Fesler

IPv6 @ Home

A few ISPs finally offer native IPv6 to consumers.

If not available: use a tunnel brokers, e. g.:

- SixXS
- Hurricane Electric

Security Consideration

- IPv6 restores end-to-end connectivity
- no NAT $\Rightarrow$ no implicit packet filter
- problem: insecure L2 net and First Hop Security (just like with IPv4)
- (still) missing support in common tools (firewalls, IDS, config UIs, etc)
- little experience

Tips for transition

- external connectivity matters
- your printer does not need IPv6
- neither does your DB backend (probably)
- start with an additional load balancer or reverse proxy

CAIDA'S IPv4 & IPv6 AS Core
AS-level INTERNET GRAPH
Archipelago January 2013

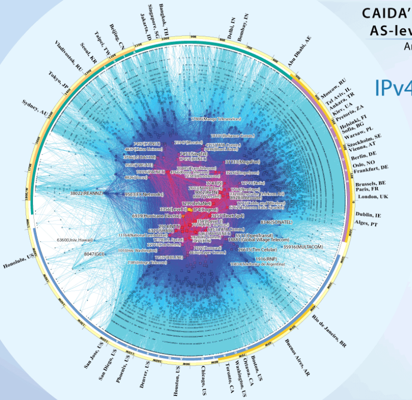

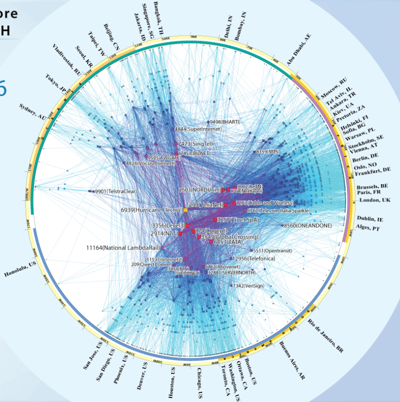

Data Source

This visualization represents a macroscopic snapshot of IPv4 and IPv6 Internet topology samples captured in 2013. The plot illustrates both the extensive geographical scope as well as rich interconnectivity of nodes participating in the global Internet routing system.

For the IPv4 map, CAIDA collected data from 58 monitors in 29 countries on 6 continents. Coordinated by our active measurement infrastructure, Archipelago (Ark), the monitors probed paths toward 214 million /24 networks that cover 93.5% of the routable prefixes seen in the Route Views Border Gateway Protocol (BGP) routing tables on 2 January 2013. For the IPv6 map, CAIDA collected data from 26 IPv6-connected Ark monitors located in 18 countries on 4 continents. This subset of monitors probed paths toward 2 million IPv6 addresses, which represent 82.2% of the globally routed IPv6 prefixes seen in Route Views BGP tables on 2 January 2013.

We aggregated this IP-level data to construct IPv4 and IPv6 Internet connectivity graphs at the Autonomous System (AS) level. Each AS approximately corresponds to an Internet Service Provider (ISP). We map each observed IP address to the AS responsible for routing traffic to it, i.e., to the origin (end-of-path) AS for the IP prefix representing the best match for this address in BGP routing tables collected from Route Views.

The position of each AS node is plotted in polar coordinates (radius, angle) calculated as indicated in Figure 1.

$$\text{radius} = 1 - \log\left(\frac{\text{transit.degree(AS)}+1}{\text{maximum.transit.degree}+1}\right)$$

$$\text{angle} = \left(\begin{array}{c}\text{longitude of the AS's}\\ \text{BGP prefixes in netacq}\end{array}\right)$$

Figure 1. Coordinates of AS in AS core.

Changes in the graph since 2012

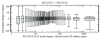

Figure 2: Between 2012 and 2013, a similar number of IPv4 ASes had degree increase and decreases. The drops were usually larger than the jumps, except for ASes with moderate degrees (between 35 and 150) which had larger jumps than drops. Among the 25 most highly connected ASes, Terremark (AS 23148) and VW Fiber (AS 19151) dropped their degree by over 92%.

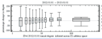

Figure 3: Between 2012 and 2013, for the 1,924 IPv6 ASes with degrees less then 4, the same number of ASes had degree increases as decreases. For larger-degree ASes, a slight majority increased their transit degree since 2012.

As in previous years, our IPv6 graph saw greater relative growth then IPv4, with 26% more ASes and 17% more AS links. (In IPv4 the growth was 11% more ASes and 6% links). These growth numbers hide a great deal of churn. Figure 13 shows that for the ~2K IPv6 ASes with degrees less then 4, about the same number of ASes increased as those that decreased their degree. The half of ASes represented by the whiskers outside the boxes increased or decreased their degree by over 30% between 2012 and 2013. Over the last year our IPv6 graph lost 295 (15%) ASes and almost 3K (38%) links, but gained 790 (41%) new ASes and over 4K (55%) new links. Our IPv4 graph was more stable, it lost ~2K (7%) ASes and ~38K (37%) links, but gained ~5K (18%) new ASes and ~43K (42%) new links. The net change in number of ASes was 495 (+26%) in our IPv6 graph and ~3K (+11%) in our IPv4 graph.

In both our IPv4 and IPv6 graphs, small and large (degree less than 5 or greater than 100) ASes split evenly between those that increased and decreased their degrees from 2012 to 2013. Drops in the IPv4 graph tended to be larger than increases; and increases in the IPv6 graph tended to be larger than drops. This reflects the faster growth of the IPv6 topology, especially for larger transit ASes.

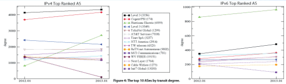

Figure 4: The top 10 ASes by transit degree.

Figure 4 plots the top 10 ASes by transit degree in either 2012 or 2013, which includes a set of 11 ASes in IPv6 and 12 ASes in IPv4. Half of these 12 most highly connected ASes in IPv4 increased and half decreased their transit degree from 2012, with a range from +234% for Hurricane Electric (AS 6939) to -92% for WV Fiber (AS 19151). Ten of the 11 top IPv6 ASes increased their degree since 2012. One AS, Init7's (AS 13030), had their transit degree drop from 207 to 91 (56%), but Hurricane Electric (AS 3356) and NTT (AS 2914) grew their transit degree over 30% to 965 and 310 respectively.

Although the set of ASes with the largest transit degrees in both IPv6 and IPv4 are increasingly converging, major differences remain. Hurricane Electric (AS 6939), the AS with the largest transit degree in IPv6, has a degree 192% larger than Level 3 (AS 3356), the second largest IPv6 AS by degree. But despite Hurricane's huge increase in IPv4 transit degree between 2012 and 2013, Level 3 (AS 3356) still has a transit degree 159% larger than Hurricane Electric's (AS 6939), and only 4% larger than Cogent (AS 174), the second largest AS by transit degree in IPv4.

ANALYSIS TEAM Bradley Huffaker, kc claffy
SOFTWARE DEVELOPMENT Young Hyun, Matthew Luckie
POSTER DESIGN Justin Cheng

	Number of IP addresses	Number of IP links	Number of ASes	Number of AS links
IPv4	27,954,132	23,494,835	34,082	109,354
IPv6	36,055	91,420	2,419	8,881

ARK HOSTS AARNet, AFRINIC, AMS-IX, APAN, ARIN, ASTI, Azrea, CAIDA, CENIC, CNNIC, CNRST, Canarie, Colorado State University, DCSI Pte Ltd, DePaul University, Evolva Telecom, Eötvös Loránd University (ELTE), Foundation for Research and Technology - Hellas (FORTH), FunkFeuer, HEAnet, Hurricane Electric, Indonesian IPv6 Task Force, International Computer Science Institute (ICSI), Internet Systems Consortium, Iowa State University, Jaguar Network, KREONet2, Kantonsschule Zug, Level 3 Communications, Men and Mice, NCAR, NIC Chile, NIC Mexico, NORDUnet, NTT, Northeastern University, Ottawa Internet Exchange, Public University of Navarra, Purdue University, QCell, RIPE NCC, RNP, Rede ANSP / Projeto NARA, Registro.br, SURFnet, Simula Research Laboratory, Southern Methodist University, TKK, TWAREN, Tinet, TorIX, UCAID, US Army Research Lab, Universitat Leipzig, Universitat Politecnica de Catalunya, University of Cambridge, University of Hawaii, University of Limerick, University of Melbourne, University of Napoli, University of Nevada at Reno, University of Oregon, University of Waikato, University of Washington, University of Zurich

COOPERATIVE ASSOCIATION FOR INTERNET DATA ANALYSIS
San Diego Supercomputer Center . University of California, San Diego
9500 Gilman Drive, mc0505 . La Jolla, CA 92093-0505 . 858-534-5000
http://www.caida.org/research/topology/as_core_network/