

# **It's 2008 – Where's My Flying Car, Err, IPv6?**

Kevin Lahey  
lahey@isi.edu

# The Basic Problem

129.99.23.29

4 octets =  $4 \times 8 = 32$  bits

4,294,967,296 addresses, give or take

Well, give or take about a billion, so it's really only about 3 billion addresses

How many people are there in the world? How many want IP addresses?

# Who Are We Talking About?

- IETF – Internet Engineering Task Force
- Working Group – subset of the IETF
- IANA – Internet Assigned Numbers Authority
- RFC – Request For Comments

# IPv6 Timeline

- 1992 – Oh, my, we might have a problem
- 1995 – Here's a spec, RFC1883
- 1996 – Prototype implementations, 6bone
- 1999 – Common OSes provide IPv6 as part of releases
- 2002 – I give up completely on IPv6
- 2003 – Windows 2003 released with production IPv6 stack
- 2008 – ???
- 2011 – No more IPv4 addresses?

# Not To Play Chicken Little...

Over the last year, the Internet community has come to realize that the IANA and RIR free pool of IPv4 addresses will be exhausted within 3-4 years, and possibly sooner. For example, Geoff Huston's "IPv4 Address Report" strongly suggests that IANA free pool will be exhausted by summer 2010, with the remaining RIR pool exhausted by summer 2011.

-- draft-narten-ipv6-statement-00.txt

# What Else Might Happen

- A new black market in IPv4 addresses?
- An officially sanctioned market in IPv4 addresses?
- NATs everywhere?
- You have a clever idea and make a fortune... (IPsec vs. SSH, anyone?)

# What Are We Talking About Here?

- Internet Protocol, building block of the Internet
- Provides best-effort packet delivery for upper layer protocols, like TCP and UDP
- TCP, of course, provides the basis for HTTP and SSH and all that good stuff
- UDP is used by DNS, NFS, TFTP, etc.
- IP takes care of moving packets over lower layers, like Ethernet or WiFi.
- Currently running IPv4.

# Why IPv6?

- Larger addresses. Duh.
- Autoconfiguration\*
- Mandatory security\*
- Generally cleaner design should allow for faster, more efficient routing (sorta)
- We are more careful handing out addresses this time, to shrink the size of the default free zone?

\* Ummm, well, yeah, IPv4 does that now, too...

# While We're At It, What Else Can We Hack?

- Make it easier to process headers, by aligning them on 64-bit boundaries.
- Support larger MTUs, up to 4 GiB?
- Deal with oversights in IPv4, like small fragmentation fields?
- In fact, get rid of in-route fragmentation altogether, and require Path MTU Discovery
- Bump up the minimum MTU, while we're at it

# Addressing

- We have 32 bits in IPv4.
- Why not go whole hog? 128 bits!
- This provides plenty of room for expansion and allows for autoconfiguration and other cool extensions.
- This has some obvious storage issues -- I hope that you have four times as much space to store routing tables on your expensive router!

# Think 129.99.23.29 Was Tough to Remember?

- To write 128-bit IPv6 addresses, use hex for every two octets, separated by colons:

2001:0db8:0000:0000:021e:37ff:f31a:4a8b

- Omit leading zeros, abbreviate a long string of zeros as '::'

2001:db8::21e:37ff:f31a:4a8b

- Ummm, yeah, you might wanna start using DNS

# Reserved Address Ranges

- Link-local addresses are available for nodes on the same physical subnet – fe80::/10
- Unique local addresses provide RFC1918-like addresses (10/8, 172.16/12, 192.168/16) for IPv6, but with some funky complications – fc00::/7
- Multicast – ff00::/8
- Localhost – ::1





# What Changed?

- Fixed header length
- No more fragmentation? Well, sorta
- Time to Live became Hop Count? Well, that's just truth in advertising.
- No more header checksum?
- What about all that jumbogram hoo-hah?  
The payload length is still 16 bits!



# Fun(?) Fragmentation Facts

- IPv4 minimum MTU – 68 octets(!?)
- IPv6 minimum MTU – 1280 octets
- What about ATM, with 48 octet payloads?  
Same as IPv4, you have to provide a mapping to make it appear to the upper-layer protocol (IPv6) that it successfully carries larger packets
- Hopefully, it'll be easier on routers to avoid fragmentation
- Plus, those 16-bit fields in IPv4 are getting too small to prevent accidental misassembly



# Option Format

```
+-----+
| Next Header | Hdr Ext Len |
+-----+
|
.
.
.
|
+-----+
```

-- RFC2460

# Jumbogram

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         | Option Type | Opt Data Len |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                         | Jumbo Payload Length |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
```

-- RFC2675



# Neighbor Discovery

- No more ARP
- Neighbor discovery is performed using ICMP
- Assemble multicast destination based on IPv6 destination address
- Interfaces listening for these multicast packets so they can respond
- No more broadcast storms (?)





# Autoconfiguration

- We want to be able to assign new IPv6 addresses, and bring up host in new environments quickly
- Do a Router Solicitation to query for routers
- Listen for a Router Advertisement to tell us prefixes and default routes and all sorts of stuff (MTU?) for the subnet
- Host can then combine MAC address (EUI-64) and prefix to determine a complete IPv6 address
- Eh, or you can use DHCPv6





# DNS

- Okay, you can't fit 128-bit addresses into regular A records
- Gonna have to use the new AAAA records
- Wow, there can be some ugliness when you get a nice AAAA answer, and don't have IPv6 connectivity (more later)

# APIs

- So, great, now I just need to install it on my OS and I'm good to go, right?
- Not so much – every C program that uses the sockets API will have to be tweaked to allow for the new larger addresses
- You also have to deal with the AAAA addresses
- Use the new `getaddrinfo(3)` instead of `gethostbyname(3)`, et al

# The Good News

- Many, many of the common applications are already good to go
- Web browsers, mail servers, mail readers, ssh, lots of other tools are (by now) ready out of the box.
- This sure could be a pain for locally custom applications

# How Do I Get IPv6?

- Hurricane Electric tunnel broker
- Vendors (more and more)
- Teredo(?)

# What Should Scare You?

- IPv6 isn't nearly as well tested as IPv4 (ummm, you can fix this)
- Who knows what new security holes are lurking out there (RH0, anyone?)
- Yow, suddenly you need four times as much storage for routing tables, and maybe IPv6 isn't going to do much to reduce the size of the default free zone
- You might also need four times as long to move those addresses around – what happens to router performance!?