COMP312-09A Communications and Systems Software

Lecture 12 – IPv6

Matthew Luckie
mluckie@cs.waikato.ac.nz

Motivation

- IPv4 is going to run out of addresses soon
  - In the early 1990s, it was clear even then
- IPng IETF working group formed, to develop a next generation Internet Protocol.
  - More addresses was the main goal
  - Decided to start with a clean slate
    - Simplify IP to make routing faster
    - Reduce routing table size
- IPv6 emerged
  - Not backwards compatible with IPv4
  - i.e. two independent IP protocols, independent routing.

What is wrong with IPv4?

- 32-bit address space is inefficiently allocated
  - Class A = 127 networks of 16 million addresses
  - Class B = 16384 networks of 65535 addresses
  - Class C = 2 million networks of 255 addresses
- Class B address space is close to exhaustion
  - How many networks can you think of with 16 million hosts?
  - How many networks can you think of with more than 255 hosts?

Demand for addresses

- Mobile phones
- Game consoles
- Light switches
- Toasters

IPv4 Header

0 15 16 31

<table>
<tr><td>Ver</td><td>HLEN</td><td>DSCP</td><td>ECT</td><td>CE</td><td colspan="3">Total Length</td></tr>
<tr><td colspan="5">Identifier</td><td>DF</td><td>MF</td><td>Fragment offset</td></tr>
<tr><td colspan="2">TTL</td><td colspan="3">protocol</td><td colspan="3">Header Checksum</td></tr>
<tr><td colspan="8">Source IP Address</td></tr>
<tr><td colspan="8">Destination IP Address</td></tr>
</table>

What is wrong with IPv4?

In 1993, it was claimed[1] that the routing tables in core routers could be reduced from 10,000 entries to 200 if addresses were re-allocated according to continental boundaries and service providers

1. Ford, Rekhter, Braun (1993). Improving the Routing and Addressing of IP.

in Stevens, W.R. (1994). TCP/IP Illustrated, Volume 1: The Protocols.

Note: BGP routing tables in 2009 have 290,000 entries

What is wrong with IPv4?

- Fragmentation is difficult for a router to do
- Until recently, TOS had no real purpose
- A 576 byte MTU is applicable to X.25 networks
  - not applicable any more
- The IP Header's checksum is wasteful for a protocol that operates on a 'best-effort' basis

Version	Traffic Class	ECT	CE	Flow Label
Payload Length			Next Header	Hop Limit
Source Address				
Destination Address				

IPv4 and IPv6

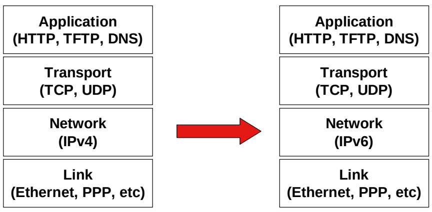

Theory goes that we just have to change the network layer to reach utopia.

IPv6: IP Next Generation

- 128 bit addresses. 6 x 10^{23} unique addresses per square metre of the surface of the earth
- No more checksum in the IP header. It wasn't protecting much anyway
  - Payload is still protected by transport-layer checksum
- Minimum MTU is now 1280 bytes.
- Fragmentation handled by end-hosts only.
  - hosts do path MTU discovery, or they send packets no larger than 1280 bytes

Representing an IPv6 address

- A 128 bit address is four times larger, in terms of space, than a 32 bit address.
- Hard to remember numeric addresses.
  - Even more dependence on DNS than before
- Textual representation is not a dotted quad like in IPv4
- Example:

  `2001:0468:0901:000a:021b:21ff:fe1e:8c5d`

  Eight four-hex numbers, each separated by a colon

Representing an IPv6 address

`2001:0468:0901:000a:021b:21ff:fe1e:8c5d`

Eight four-hex numbers, each separated by a colon

`2001:468:901:a:21b:21ff:fe1e:8c5d`

Can remove leading zeroes in each quad

Representing an IPv6 address

`2001:0dc0:2001:0000:4608:0020:0000:0000`

Eight four-hex numbers, each separated by a colon

`2001:dc0:2001:0:4608:20:0:0`

Remove leading zeroes in each quad (need to keep one zero in each quad though)

`2001:dc0:2001:0:4608:20::`

A single string of 8 zero digits can be replaced with ::

IPv6 address allocation

- The IPv6 address space is immense.
- IANA is delegating to regional registries out of a /4 prefix.
  - i.e. 1/8th of the address space. If we use the address space poorly, then there is still 7/8ths left.
- APNIC minimum initial IPv6 allocation is a /32
  - i.e. ISP receives 96-bits, and can delegate it how it pleases.
  - About 20 000 000 000 000 000 000 times the size of total IPv4 space
  - Registry also reserves prefixes around initial allocation, so it is possible for an ISP to be delegated a /30 or /28 without having to renumber their network

IPv6 address allocation

- Addressing model says each customer of an ISP should ideally be assigned a /56
  - i.e. bits 32 to 56 (24 bits, or enough for 16,777,216 customers)
- Each customer should use /64 prefix
  - i.e. bits 56 to 64 (8 bits, or 255 subnets)
  - A single /64 provides 2^{32} times the total IPv4 address space
- Clearly some immense waste apparent
  - Hopefully the immenseness of the address space is larger than the immenseness of the waste.
  - A model only, ISPs and customers can use their addresses however they like in reality

IPv6 addresses in the DNS

- IPv6 addresses have a separate record type in the DNS
  - AAAA
- Client name resolvers now ask for
  - AAAA (IPv6 address records) and
  - A (IPv4 address records) when a name is looked up
- Most resolvers return the AAAA records first, meaning client applications will prefer IPv6 addresses over IPv4 addresses

Java Applications

- Java applications should work with IPv6 without any extra programming required

`Socket(InetAddress address, int port)`

InetAddress can hold both IPv4 and IPv6 addresses, transparently to the application programmer.

Other programming languages and systems aren't so good.

IPv6: IP Next Generation

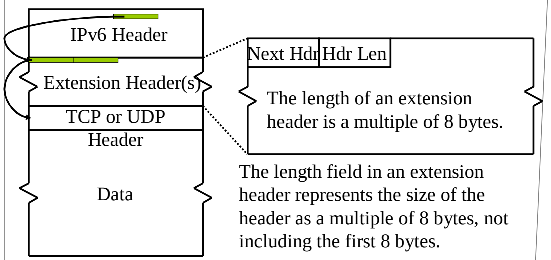

The length field in an extension header represents the size of the header as a multiple of 8 bytes, not including the first 8 bytes.

IPv6 Extension Headers can be chained together in a linked-list.

IPv6: IP Next Generation

- IP header is a constant size (can tell this by no HLEN field in the header)
- IPv6 extension headers can be thought of as IP options.
- IPv6 Extension Headers add functionality
  - hop-by-hop options
  - packet fragmentation
  - authentication and encryption (IPSec)

Hop-by-Hop options

- The Hop-by-Hop options format is Type-Length-Value (TLV)

Option Type	Opt Data Length	Option Data

Hop-by-Hop options

- IPv6 packets are limited to 64KB in size with a 16-bit IPv6 header length field. This can be increased to 4GB with the jumbogram option.
- When a Jumbogram is used, the length field in the IPv6 header is set to zero.

		Type = 194	Length = 4
Jumbo Payload Length			

Packet Fragmentation

Next Header	Reserved	Fragment Offset	R	M
Identification				

R: Reserved
M Flag: 1 = more fragments
0 = last fragment

Packet Fragmentation can ONLY be performed by end-hosts, not by routers

Migration Path

- Native IPv6 is the goal
- Until then, we're going to be using some transition mechanism between IPv4 and IPv6
- I.e. we're going to have pockets of IPv6 around the Internet, and we want to connect them together

IPv6 over IPv4 tunnelling

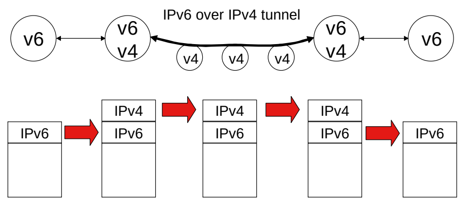

Where are we now? IPv6 routing table

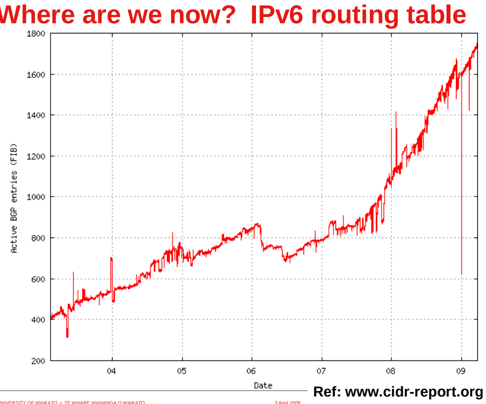

Ref: www.cidr-report.org

Where are we now? IPv4 routing table

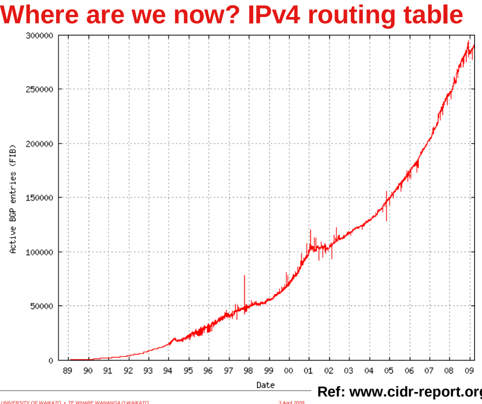

Ref: www.cidr-report.org

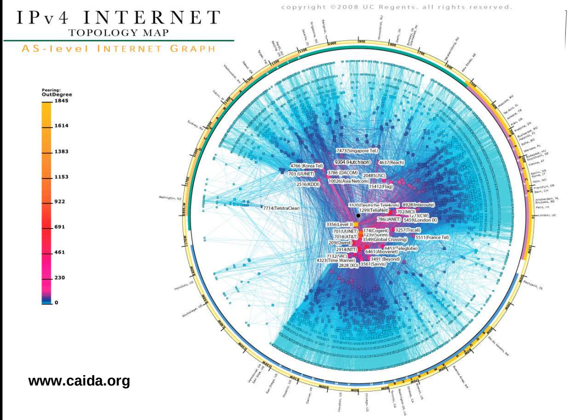

www.caida.org

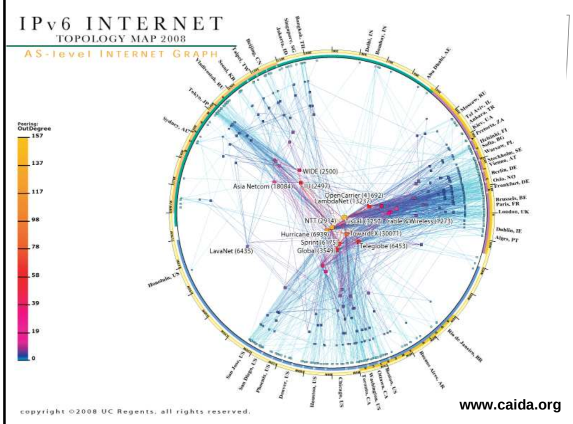

www.caida.org

IPv6 support

- Most (not all) networking vendors have IPv6 support in their products
- Most operating systems support IPv6
  - FreeBSD
  - Windows 2000, XP, Vista
  - MacOS X
  - Linux
  - Quality of implementations and usefulness to programmer has improved over time.
- Not all applications easily support IPv6
- Vista's IPv6 stack is on by default, XP's is not.
  - Vista's support for IPv6 in applications is much improved over XP's.

Summary

- IPv6 defined in 1998, still seeing limited deployment
- Some networks in NZ can route IPv6
  - FX
  - TelstraClear
- IPv4 will probably run out soon. We are not sure what effect, if any, this is going to have on IPv6 deployment.