COMP312-09A Communications and Systems Software

Lecture 11

Address Exhaustion, RFC 1918, Network Address Translation

Matthew Luckie

mluckie@cs.waikato.ac.nz

THE UNIVERSITY OF WAIKATO

IPv4 Header

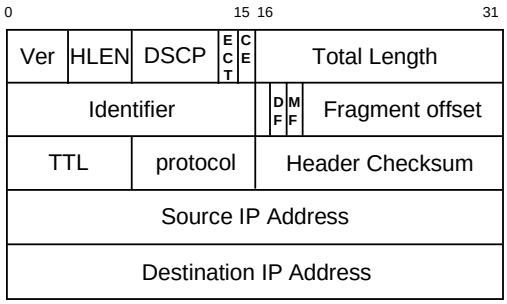

IPv4 Address Exhaustion

- A 32 bit address allows for around **4 billion** hosts, ignoring network and broadcast addresses
  - 2^{32} = 4,294,967,296
- This might seem a lot, until you think about giving an IPv4 address to all cell phones in the world
  - This scenario was not envisaged when IPv4 was defined
- Ignoring this, class-based assignment resulted in inefficient allocation of IPv4 address space

Two intermediate solutions

- Classless Inter-domain Routing (CIDR)
  - RFC 1519, September 1993
- RFC 1918 private addresses
  - originally RFC 1594, March 1994
  - Requires application level gateways, or
  - Network address translation (RFC 1631, May 1994)
- Both bought the Internet time before address run out
- Long term solution
  - IP version 6 – next lecture

IPv4 Exhaustion Counter

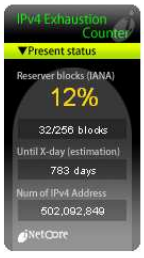

This gadget has been developed by Takashi Arano's Intec NetCore.
Details at http://entne.jp/tool/toollist/000101.php

IANA allocations

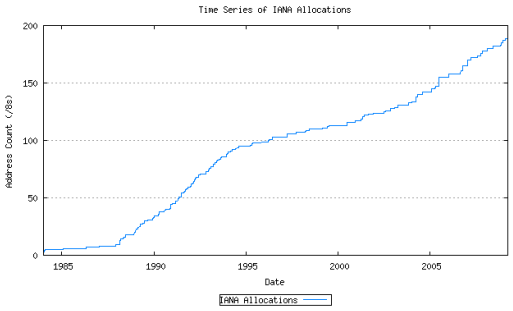

source: http://ipv4.potaroo.net

IPv4 Address Status

source: http://ipv4.potaroo.net

Classless Interdomain Routing (CIDR)

- Up until 1993, allocations were done in classes
  - Class A: /8
  - Class B: /16
  - Class C: /24
- Changes to allocation policy and routing protocols allowed allocations to be made based on actual space required.
  - E.g. allocating 202.53.176.0/20 to FX networks, 4096 addresses
    - Instead of allocating one class B network, for 6% usage
    - Instead of allocating 16 class C networks

RFC 1918 private addresses: motivation

With the proliferation of TCP/IP technology worldwide, including outside the Internet itself, an increasing number of non-connected enterprises use this technology and its addressing capabilities for sole intra-enterprise communications, without any intention to ever directly connect to other enterprises or the Internet itself.

RFC 1918 private address space

- 10.0.0.0/8
  - Class A network
- 172.16.0.0/12
  - Set of 16 contiguous class B networks
- 192.168.0.0/16
  - Set of 256 contiguous class C networks

- Private addresses are, by definition, available for anyone to use for their own networks
- They do not uniquely identify any host in the Internet
- They are not routed on the Internet

Internet End-to-End Connectivity

- Ideally, all hosts in the Internet would be able to uniquely identified by their address
  - Allows for simple inter-connectivity of hosts
- The use of RFC 1918 addresses breaks this model

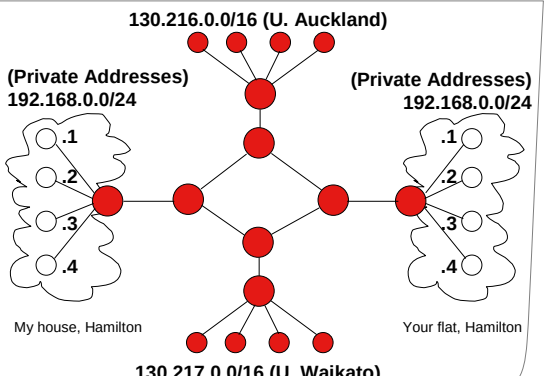

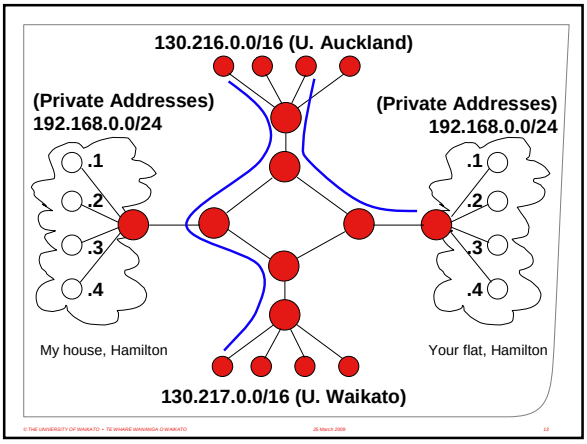

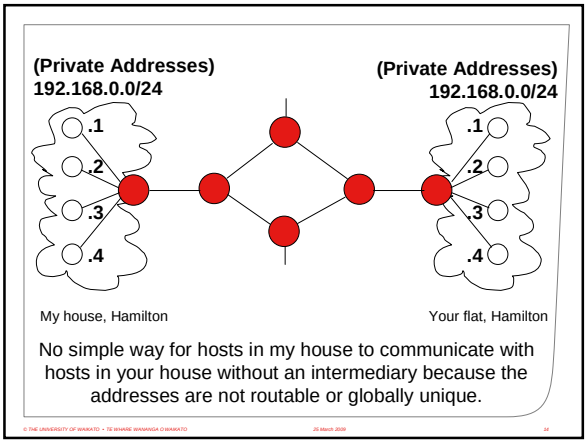

No simple way for hosts in my house to communicate with hosts in your house without an intermediary because the addresses are not routable or globally unique.

Inbound connections

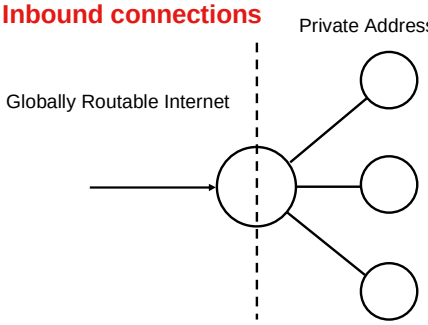

Not at all obvious how we should forward inbound connection requests – which host should it go to?

Workarounds for RFC 1918

- Application Level Gateways (ALGs)
- Network Address Translation (NAT)
- Internet Protocol Version 6 (IPv6)

- None are perfect.

Application Level Gateways

- End host has an assigned RFC 1918 address
- Application configured to use an 'application level gateway' which will handle communication with the outside world
  - i.e. that which is globally routable
- ALG has RFC 1918 address, and a globally unique address
- Limitations:
  - Each application requires its own ALG protocol to be defined and implemented
  - Not easy to design protocols or ALGs to allow two hosts in separate RFC 1918 islands to communicate
    - i.e. inbound connection support

Application Level Gateways

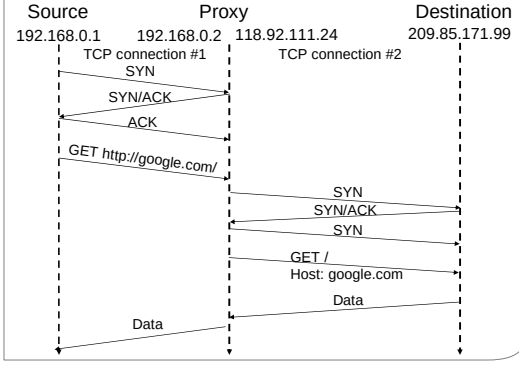

Network Address Translation

- NAT gateway configured as a router acts as transparent relay. It does this by
  - re-writing the source IP address and source port on outgoing packets
  - remembering how it translated the packet so that it can re-write the destination IP address for the reply packets

(1) <192.168.0.1, 1024> TCP maps to <118.92.111.24, 3087> for <209.85.171.99, 80>

(2) <192.168.0.8, 1024> TCP maps to <118.92.111.24, 3088> for <209.85.171.99, 80>

(3) <192.168.0.8, 53> UDP maps to <118.92.111.24, 3089> for <209.85.171.99, 53>

Network Address Translation

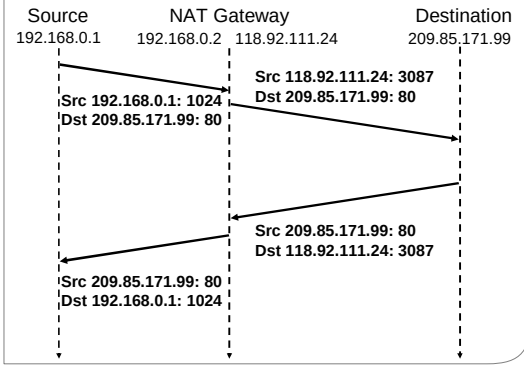

Network Address Translation

- Advantages:
  - Fairly simple to implement
  - Usable with UDP, TCP, and ICMP packets
  - Does not require separate implementation for each protocol like an ALG does
- Disadvantages:
  - Still does not allow end-to-end connectivity
    - Breaks peer-to-peer applications
    - Does not allow in-bound connections
  - Single point of failure: if NAT gateway breaks then all connections go with it, as it has state associated with each connection

IPv4 Exhaustion Counter

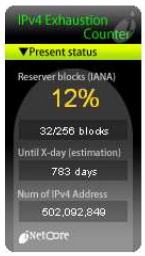

This gadget has been developed by Takashi Arano's Intec NetCore.
Details at http://entne.jp/tool/toollist/000101.php

Service Provider NAT

- In New Zealand, we are assigned a single globally routable IPv4 address whenever we connect with dial up or DSL
  - This is not true in all countries
  - i.e. customers in some countries are assigned an RFC 1918 address
- Soon, addresses will run out (within 4 years)
- The 'correct' solution is to move to IPv6 which has 2^{128} addresses
  - As far as I know, no consumer Internet service in New Zealand provides IPv6.
- If this does not happen Service Provider NAT (SPNAT) might be required

Intermission

- Shane Alcock's NZNOG 2009 SPNAT slides

Summary

- It became obvious fairly on in the early 1990s that the Internet was going to run out of addresses
  - CIDR, Private addresses bought us time
  - NAT is the price we paid
- Projected run out is within about 4 years
  - IPv6 is intended solution
  - Not at all well deployed, has been defined for over 10 years now.
  - Service provider NAT not a good solution either.
- Next lecture: IPv6

Further Reading

- Pages 441 to 444 (CIDR)
- Pages 444 to 448 (NAT)
- Pages 464 to 473 (IPv6, topic of next lecture)