MONASH University Go Boldly.

IPv6 at Monash University

John Mann

www.monash.edu

MONASH University www.monash.edu

Summary

- **IPv6 research at CTIE, AT CRC**
- **Production IPv6 dual-stack network**
- **Problems**
- **Future Plans**

MONASH University www.monash.edu

The Centre for Telecommunications and Information Engineering, and Advanced Technologies CRC

Did leading-edge research into

- **Mobile IPv6**
- **Protocols for Detecting Network Attachment**
- **Fast handovers and fast address configuration for Mobile IPv6**
- **Streaming Video over IPv6**

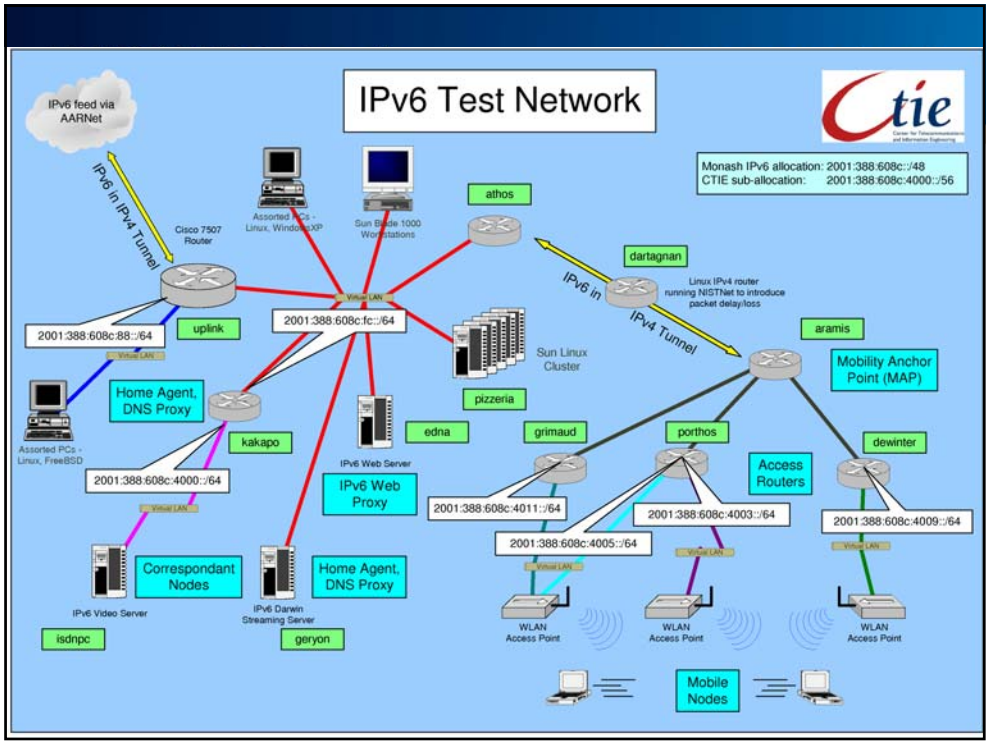

Advanced Technologies CRC

Involved
- **14 staff from Monash and RMIT**
- **50+ students**

And resulted in
- **30+ Internet Drafts**
- **RFC 4135 and RFC 4429**
- **Optimistic DAD now in Linux kernel**
- **125k lines of IPv6Suite simulation code**
- **7 PhDs**

MONASH University www.monash.edu

Production Monash IPv6 dual-stack Network

- **Why Now?**
- **Addressing Plan**
- **Configuring Routers and ACLs**
- **DNS and DHCP Services**
- **Servers**
- **Clients**

Why Now?

- **Monash is like a battleship, it will take a long time to change direction**
- **Need to break the chick-and-egg problem by providing the IPv6 infrastructure**
- **Need to raise IPv6's profile, and show that it does actually work**
- **Need to find out what we don't know**
- **Dual-stacking the infrastructure is a 50-device problem, v. dual-stacking edge devices which is a 30,000-device problem**
- **Cost of failure, or having to do things over, is low now, compared to having to do things in a rush later**
- **Be seen to be a leader**

MONASH University www.monash.edu

IPv4 address plan

- **Originally used 130.194/16 for everything**
- **When that became full, added**
  - 172.16/16 (private nets)
  - 172.17/16 (wireless)
  - 172.18/16 (students)
  - 172.19/16 (staff)
  - 172.20/16 (management)
  - ...
  - with Internet access via Proxies
- **Renumbered half the network**

IPv6 Address Plan

- No easy way to map IPv4 addresses <=> IPv6 addresses
- IPv4 address plan a bit disorganised anyway
- So, new logical IPv6 address plan

Use (2 bits)	Org Unit (4 .. 12 bits)	Location (8 .. 0 bits)
Server Research Staff Student	ITS Admin Arts …	North East West …

- /64 for p-p router backbone links
- /64 for router loopbacks, DNS anycasts

MONASH University www.monash.edu

IPv4 Address Plan 2

- In 2008, replaced Web and SOCKS proxies with Cisco Service Control Engine (SCE)
- All hosts that want Internet access needed to move to Public IPv4 addresses
- Obtained extra IPv4 addresses, and renumbered half the network again
- IPv4 network plan is now less messy – so less justification for a different IPv6 network plan

Cisco Routers

- **Catalyst 6500 routers**
  - IPv6 availability varies with releases/feature sets
  - From 12.2(33)SXI, released 13th Nov 2008,
    IPv6 is in IP SERVICES,
    previously ADVANCED IP SERVICES
- **Catalyst 3750 routers**
  - IPv6 features with ADVANCED IP SERVICES license, about $4k list price
  - IPv6 ~halves TCAM available for IPv4 functions
  - Has severely restricted IPv6 for small campuses that use 3750's as their routers

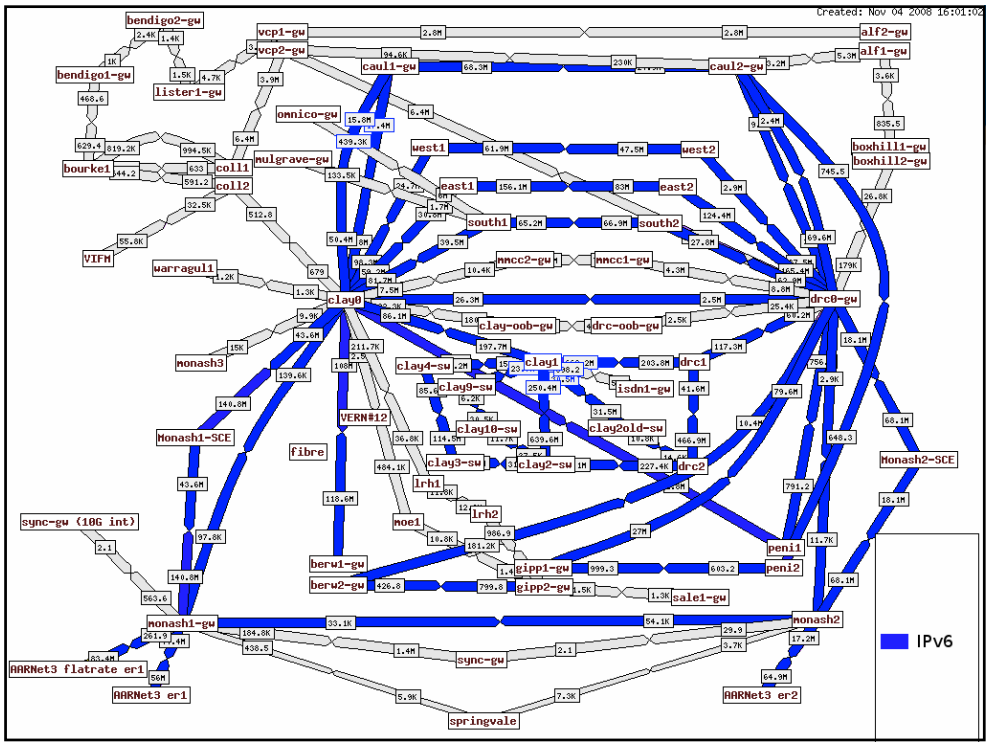

Router Configuration

- We enter all subnet information into a subnet database: subnet name, address ranges, vlan name/number, router names, access rights, …
- The database is the used to generate router configurations
- Extended the database to manage IPv6 addresses and generate IPv6 router configs

MONASH University www.monash.edu

ACL Management

- Old scheme creates IPv4 ACLs using a Perl script and a flat-file list of exceptions
- ACL heuristics in script evolved over time
- Now need to generate IPv6 ACLs as well
- New ACL creation scheme uses templates and macro expansion for both IPv4 and IPv6 ACLs
- More-transparent scheme, but lower level

IPv6 out ACL Template Example

```
permit tcp any any established
%special-top-out6
permit ipv6 %fromgroup6 any ! normal out traffic
deny udp any any range 135 139 log ! block window virus
deny udp any any eq 445 log ! block window virus
permit udp %net-monash-au6 any ! NACP
permit udp any any gt 1024 ! NACP2
permit icmp any any ! NACP
permit ipv6 any ff00::/8 ! multicast out
%special-bottom-out6
deny ipv6 any any log-input
```

MONASH University www.monash.edu

DNS

- **Addhost, our network host registration scheme, was extended to cater for**
  - Fixed IPv6 addresses; and also
  - "auto" IPv6 addresses generated from
    IPv4 address => IPv4 subnet => IPv6 subnet table,
    and Ethernet address => EUI-64 host address
- **No need for normal users to enter long hex addresses**
- **Forward and reverse DNS**
  - But not fe80::/64 reverse yet
- **DNS servers have IPv6 Anycast addresses, tied to tun0 device, advertised using Quagga**

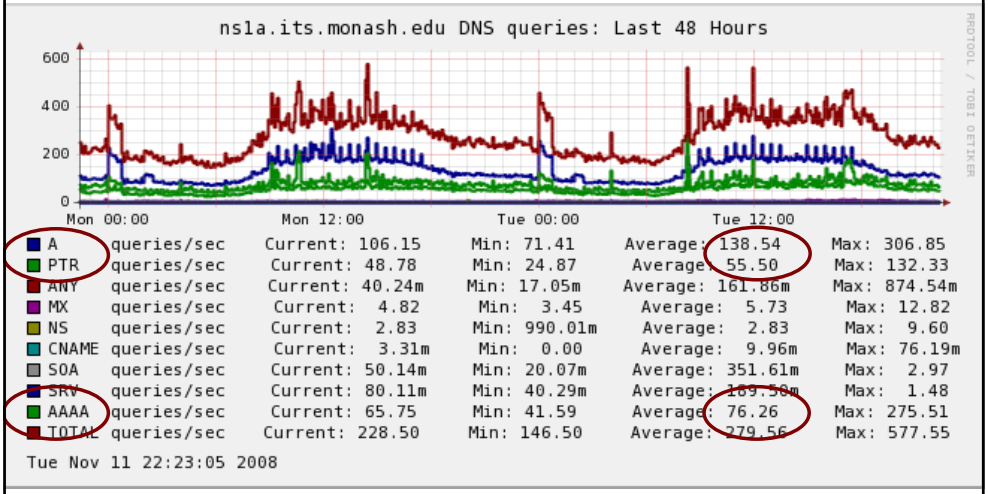

MONASH University www.monash.edu

IPv6 DNS Usage

- **About 25% of DNS queries are for AAAA records**
- **Unfortunately, most of these are for non-existent domains, or hosts without AAAA records**
- **For example: 4 AAAA queries before A query**

15:00:04.106 query: directory.monash.edu.au IN AAAA +
15:00:04.109 query: directory.monash.edu.au.its.monash.edu.au IN AAAA +
15:00:04.110 query: directory.monash.edu.au.monash.edu IN AAAA +
15:00:04.111 query: directory.monash.edu.au.monash.edu.au IN AAAA +
15:00:04.111 query: directory.monash.edu.au IN A +

- **Also occasional storms of requests for "ISATAP"**
  - 600 queries per second, per PC !!

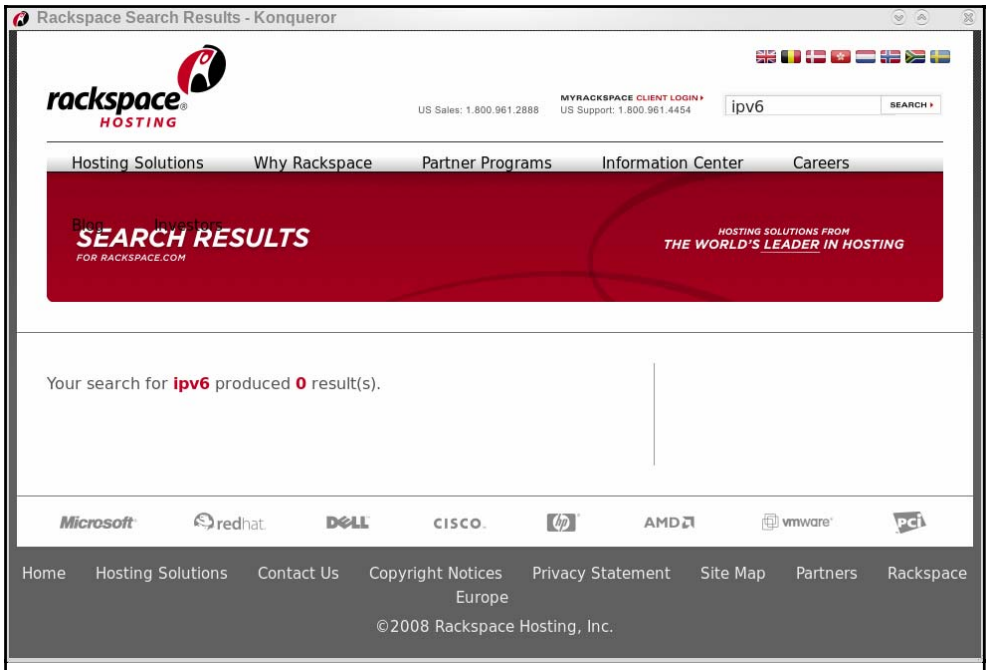

Offsite DNS Hosting FAIL

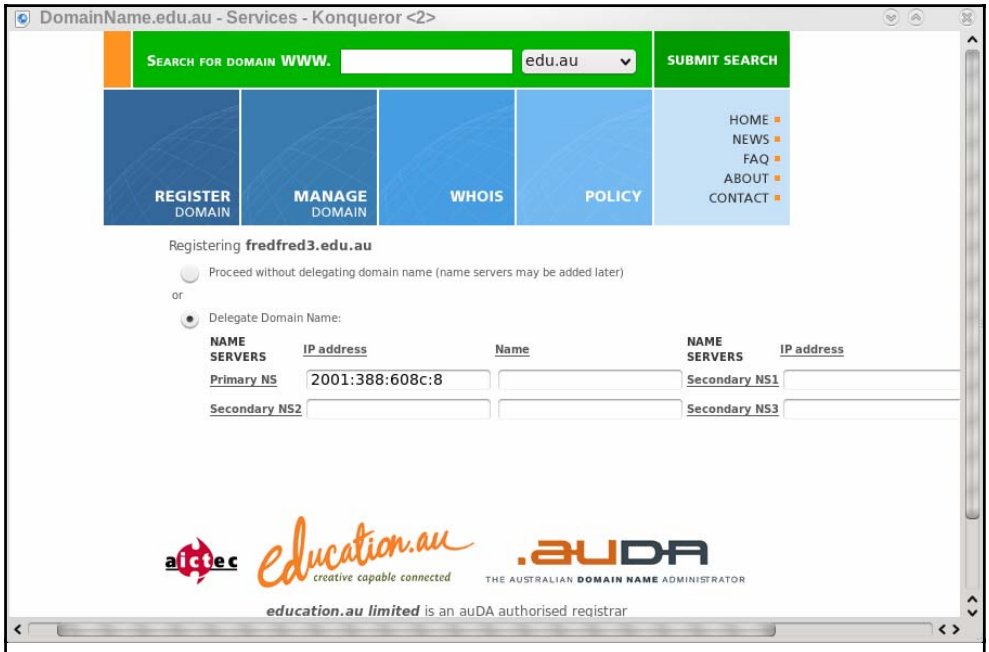

Parent DNS Zone FAIL

DHCP and network auto-discovery

- **Since we are planning to run a dual-stack network, there isn't a pressing need for IPv6 DHCP**
- **Hosts get (IPv4) DNS server addresses from IPv4 DHCP, or statically configured**
- **Network router auto-discovery has worked very well so far – have only rebooted routers once in last 1.5 years**
- **Haven't done IPv6 HSRP**
- **Beware of rogue IPv6 routers**

MONASH University www.monash.edu

Servers

- **Many servers auto-configure IPv6 by default**
- **May need to tweak**
  - /etc/hosts
  - /etc/hosts.allow
  - Ifcfg-eth0
    - > Preference to use auto-discovery IPv6 address rather than fixed IPv6 address for outgoing connections

Web Servers

- **A customised Apache 1.3 is used on our main Web serving farm**
  - Too hard to add IPv6 support
- **Apache 2.0 does support IPv6**
  - Need to check .htaccess files
    - > Permissions based on IPv4 addresses aren't relevant any more

MONASH University www.monash.edu

Web Reverse Proxy

- **Use Apache in Reverse-proxy mode as IPv6 -> IPv4 gateway**

```
Listen [2001:388:608c:88b::123]:80
<VirtualHost *:80>
   ProxyPreserveHost On
   ProxyPass / http://130.194.11.123:80/
   ProxyPassReverse /
     http://130.194.11.123:80/
</VirtualHost>
```

Diagram of IPv6 -> IPv4 Web proxy

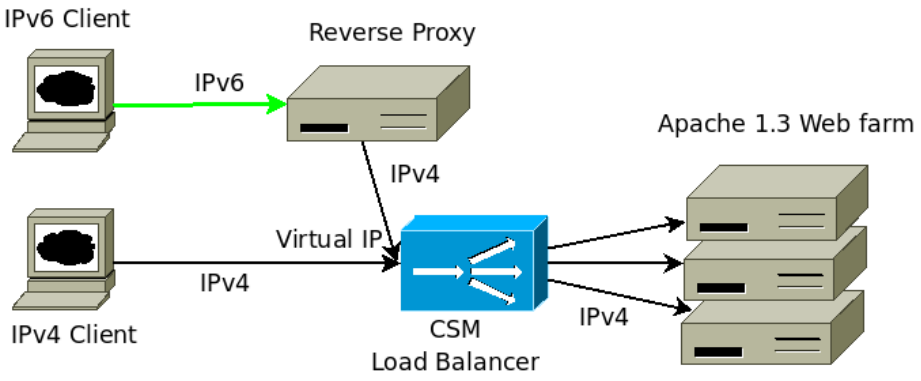

MONASH University www.monash.edu

Clients

- **Currently, end-user IPV6 subnets are limited to**
  - Electrical Engineering
  - eResearch
  - ITS
- **Biggest problem was one Vista laptop with ICS that wanted to be a IPv6 router**
- **Technical users are enjoying e.g. direct IPv6 connections between desktop at work, and PCs on home network**

Problem: Monitoring IPv6 Network

- **Statseeker V3**
  - Can monitor interface usage and up/down status
  - Can't ping IPv6 addresses
- **flow-tools**
  - Handles Cisco NetFlow V5 – IPv4 only
  - Need NetFlow V9 for IPv6
- **Fluke NetFlow Tracker 3.0.7**
  - Can accept NetFlow V9
  - Can show 6to4 IPv6 traffic
- **Snort 2.8**
  - IPv6 support is incomplete
  - Needs addresses like 2001:0:0:0:0:0:0:0/16
- **No IPv6 usage statistics collection !!!**

MONASH University www.monash.edu

Problem: IPv6 Capability of Middle-boxes

- **CSM - Content Switch Module**
  - Load-balances and routes IPv4 only
  - Could add extra IPv6-only router interface to provide IPv6 service to real server Vlan
- **FWSM - Firewall Services Module**
  - L2 mode - Pass IPv6 without inspection
  - L3 mode - IPv6 inspected, but no Multicast
- **SSL Services Module**
  - IPv4 only
- **WISM - Wireless Services Module**
  - IP protocol independent !!!
- **SCE 2000 - Service Control Engine**
  - Pass IPv6 without inspection
- **VPN 3000 Concentrators**
  - VPN over IPv4 only, IPv4 inside (could do v6-in-v4 in VPN)

Lessons Learnt

- **IPv4 not broken yet. Users see little need to migrate to IPv6, or dual-stack**
- **IPv4 address exhaustion like Y2K, or Global Warming, or Urban Sprawl, or …**
  - But no definite deadline date
  - IPv4 will continue to work after exhaustion
  - "Will effect others, not us"
- **Enabling IPv6 (over a relaxed timescale) has created opportunities for improving many IPv4 practices**
  - Address plan
  - ACLs
  - Router configuration management
  - Network monitoring

MONASH University www.monash.edu

To Do

- **IPv6 for off-site DNS secondaries**
  - University of Newcastle
  - Rackspace.com
- **AAAA records in .edu.au parent zone**
- **IPv6 E-mail services**
- **IPv6 for South Africa and Malaysia campuses**
- **IPv6 usage statistics**
- **More**
  - IPv6 servers
  - IPv6 user subnets
  - IPv6 education

MONASH University www.monash.edu

Questions?

MONASH University www.monash.edu

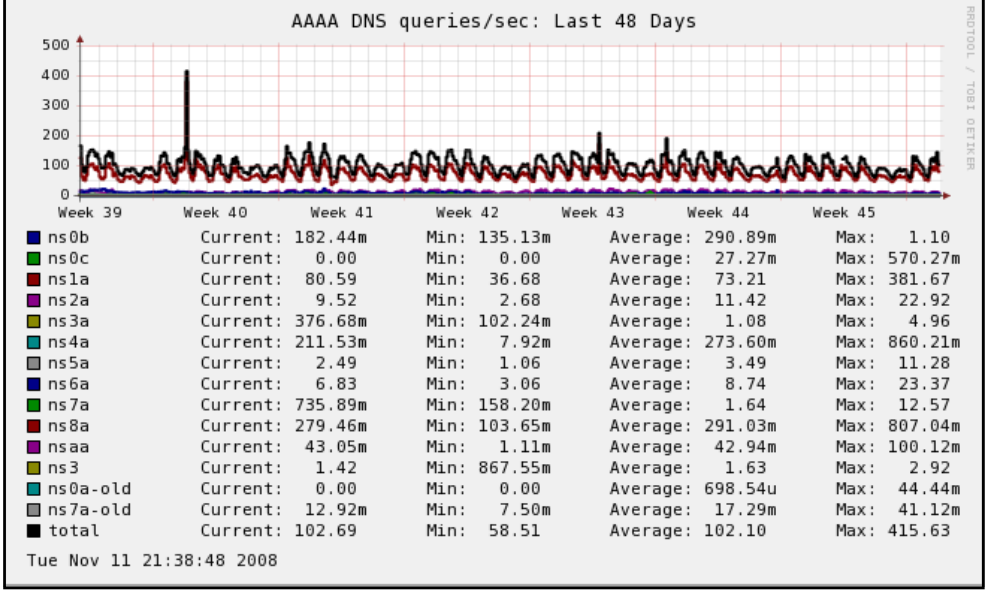

MONASH University www.monash.edu

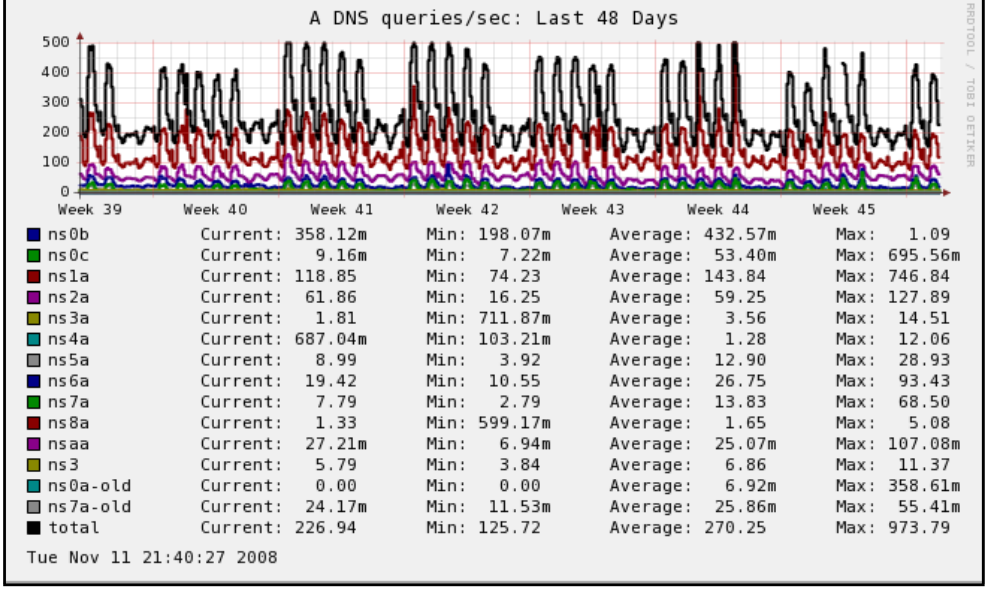

ns0b	Current: 358.12m	Min: 198.07m	Average: 432.57m	Max: 1.09
ns0c	Current: 9.16m	Min: 7.22m	Average: 53.40m	Max: 695.56m
ns1a	Current: 118.85	Min: 74.23	Average: 143.84	Max: 746.84
ns2a	Current: 61.86	Min: 16.25	Average: 59.25	Max: 127.89
ns3a	Current: 1.81	Min: 711.87m	Average: 3.56	Max: 14.51
ns4a	Current: 687.04m	Min: 103.21m	Average: 1.28	Max: 12.06
ns5a	Current: 8.99	Min: 3.92	Average: 12.90	Max: 28.93
ns6a	Current: 19.42	Min: 10.55	Average: 26.75	Max: 93.43
ns7a	Current: 7.79	Min: 2.79	Average: 13.83	Max: 68.50
ns8a	Current: 1.33	Min: 599.17m	Average: 1.65	Max: 5.08
nsaa	Current: 27.21m	Min: 6.94m	Average: 25.07m	Max: 107.08m
ns3	Current: 5.79	Min: 3.84	Average: 6.86	Max: 11.37
ns0a-old	Current: 0.00	Min: 0.00	Average: 6.92m	Max: 358.61m
ns7a-old	Current: 24.17m	Min: 11.53m	Average: 25.86m	Max: 55.41m
total	Current: 226.94	Min: 125.72	Average: 270.25	Max: 973.79

Tue Nov 11 21:40:27 2008

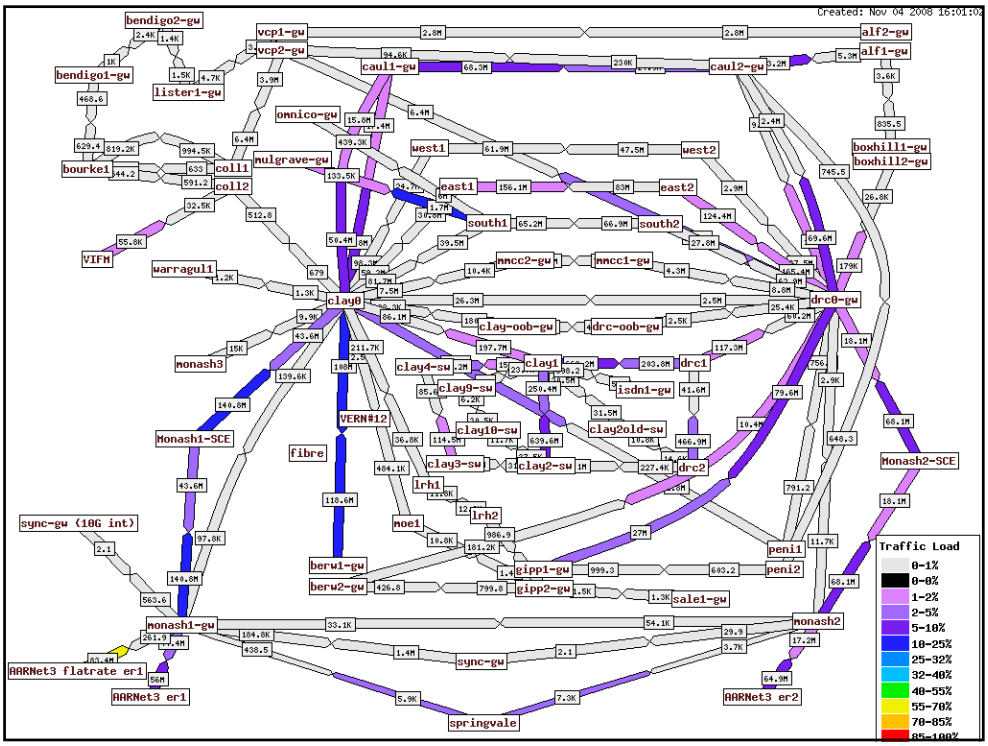

AT CRC

Major Lessons

- **IPv6 is a useful tool for 3G+ networks**
- **When funding stops, research setups and expertise is lost**

Future Work

- **Research in IPv6 and Wireless networking**
- **Emphasis is on network simulation rather than physical networking**