IPv6 Deployment at Monash University

John Mann

Agenda

- IPv6 is Coming
- IPv6 is Already Here
- Monash IPv6 Progress
- Addressing Schemes
- End Systems
- Monitoring Network Address Usage
- Traffic Monitoring
- Problems

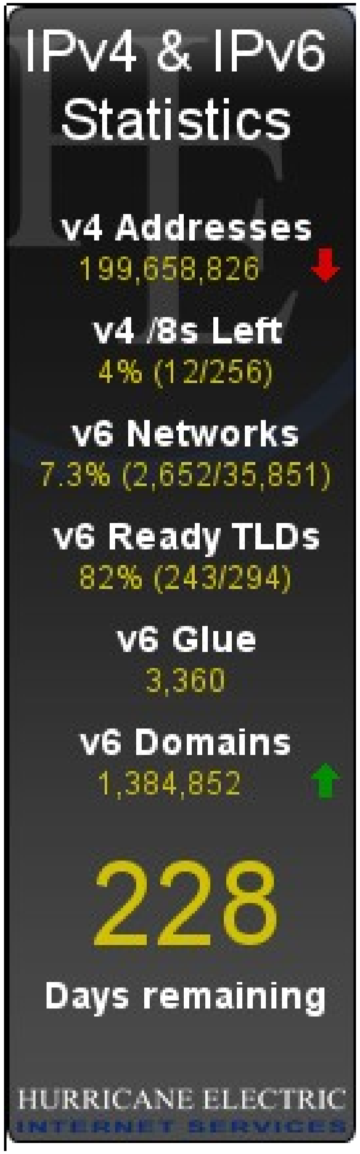

IPv4 & IPv6 Statistics

v4 Addresses
199,658,826 ↓

v4 /8s Left
4% (12/256)

v6 Networks
7.3% (2,652/35,851)

v6 Ready TLDs
82% (243/294)

v6 Glue
3,360

v6 Domains
1,384,852 ↑

228
Days remaining

HURRICANE ELECTRIC
INTERNET SERVICES

IPv6 is Coming

- **IPv6 solves the IPv4 address shortage problem by having billions of billions of more addresses**
- **IPv6 is required within 2 years**
  - Within current 3-year Strategic Plans
  - Within lifetime of equipment being bought this year
  - Within lifetime of much existing equipment
  - Within most people's current employment position
- **We will need most things to have IPv4 *and* IPv6 so they can talk to the old and new parts of the Internet, and for the new and old Internet to talk to us**
- **NAT won't solve our problems forever**
- **There is no PLAN B (IPv8 or ...)**

IPv4 Exhaustion Date has Stopped Slipping!

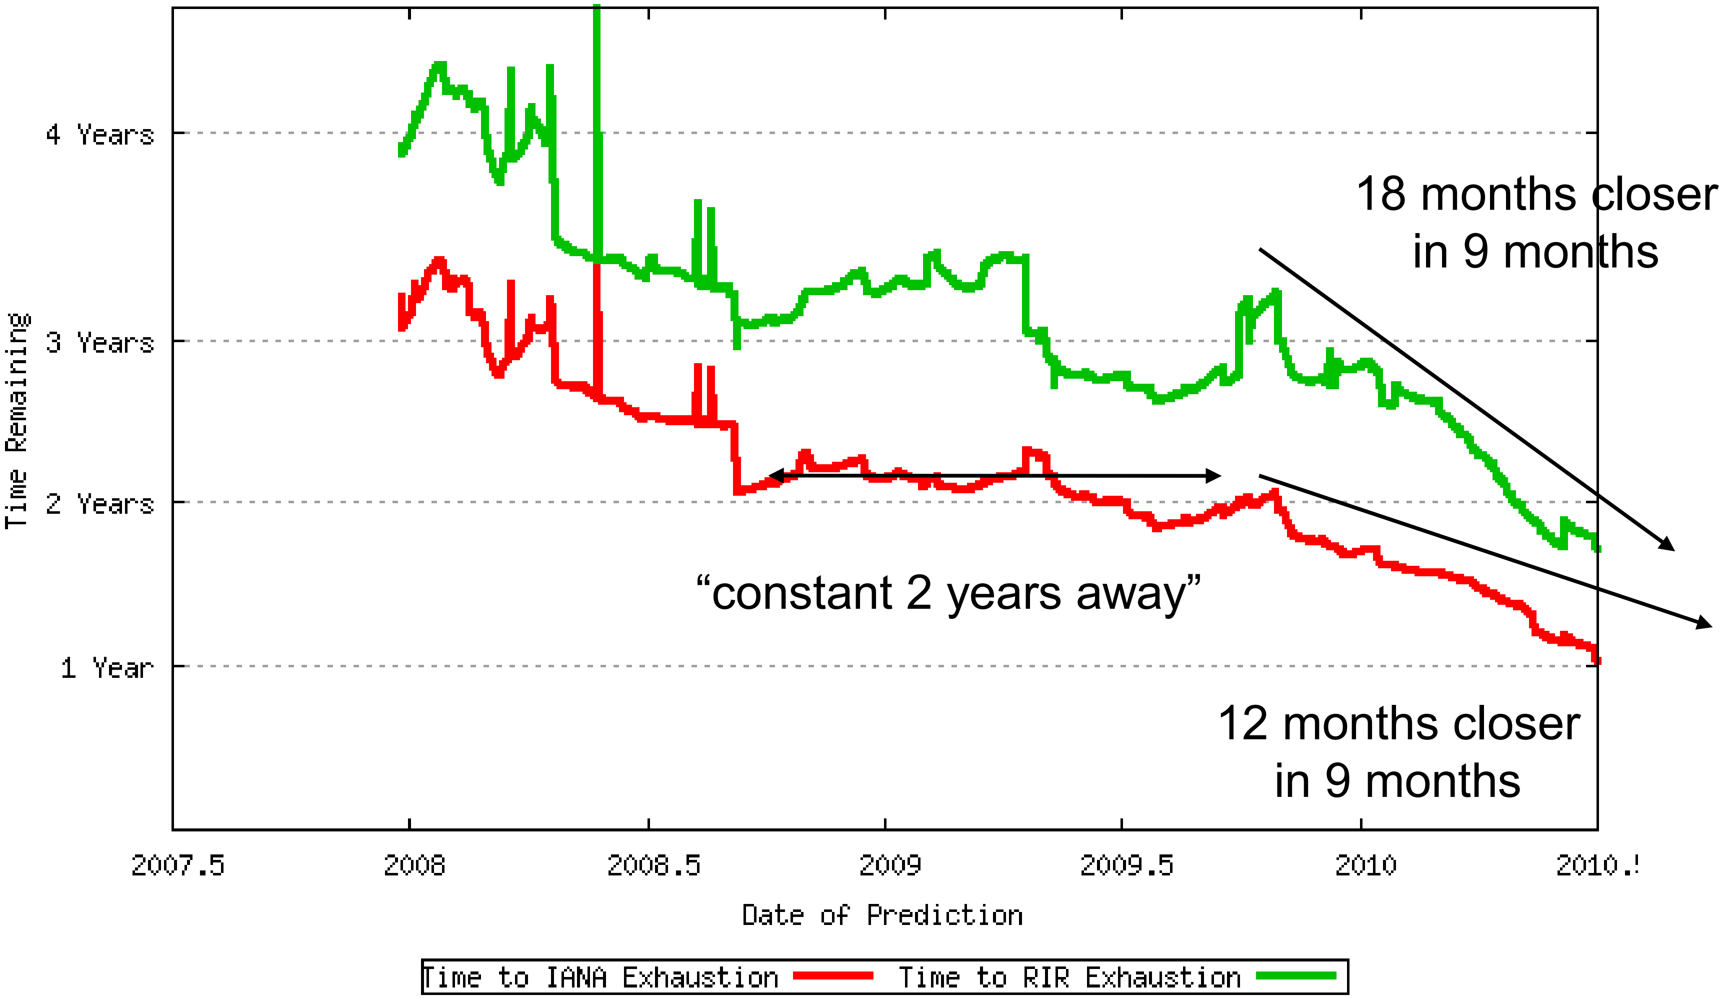

18 months closer in 9 months

“constant 2 years away”

12 months closer in 9 months

From: Geoff Huston, potaroo.net

IPv6 is already here.
It's just not evenly distributed.

- Apologies to William Gibson

- First IPv6 RFCs in 1996
- Monash has had native IPv6 since 2003
- Windows Vista / 7, Mac OS X, Linux already come with IPv6 enabled
- Some Web sites are now IPv6 enabled
  - http://www.google.com.au 23-Jan-2009
  - http://www.youtube.com/ 29-Jan-2010
  - http://ipv6.beijing2008.cn/en
- ISPs like AARNet, Internode, Vocus, NTT, HE
- Also IPv6 traffic invisibly tunneled over IPv4

Why Does Monash Need IPv6

- **Low-cost Business Continuity insurance**
  - Not costly to ensure that when the world wants to talk IPv6 to Monash, we will be ready for them
- **Monash will likely need IPv6 capability to communicate with India/China/Korea/Japan**
  - Exchange partners
  - Potential students
- **Africa (where we have a campus) may skip IPv4 and go IPv6**
- **Native IPv6 has reduced complexity**
  - Reduces cost
  - Improves network management
- **Take advantage of any IPv6 opportunities**
- **Be seen to be a leader**
- **Interesting project to keep tech experts busy**

Early IPv6 at Monash

- **In early 2000's, 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
- **Initially 6Bone tunnel via Trumpet Software**
- **Later native IPv6 over GrangeNet routed into separate research Vlans**
- **Early hand-edited IPv6 DNS**

IPv6 Test Network

Monash IPv6 allocation: 2001:388:608c::/48
CTIE sub-allocation: 2001:388:608c:4000::/56

2003

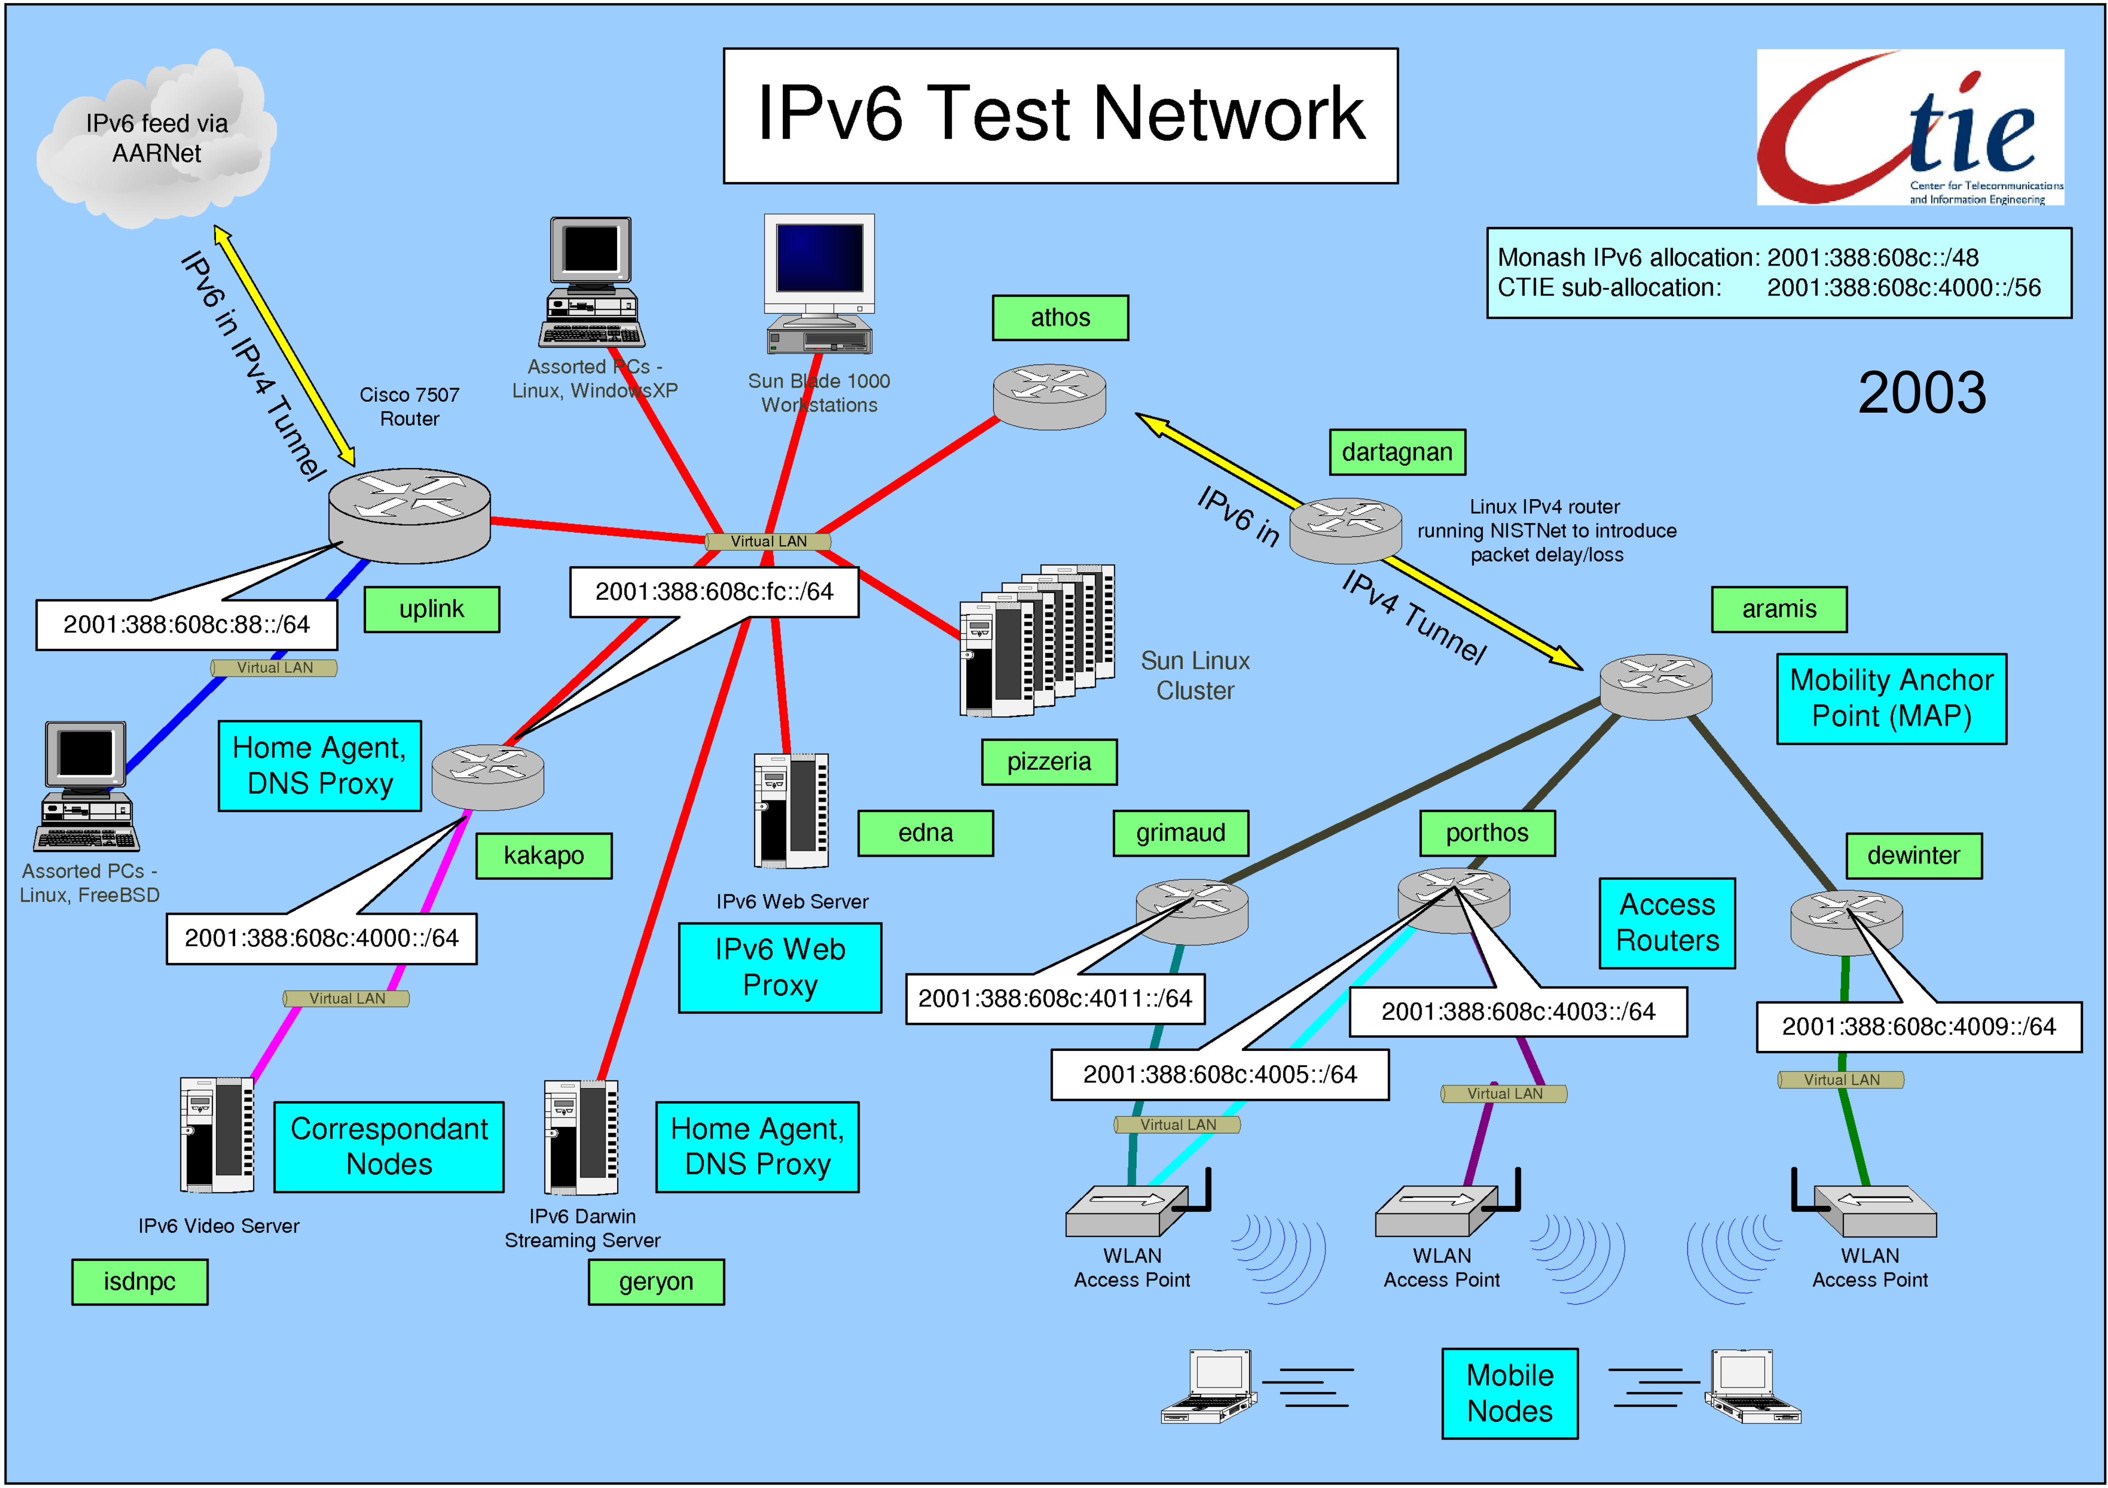

Monash IPv6 Progress So Far

- **Routers at all Victorian campuses have IPv6**
- **Services like Addhost and DNS support IPv6**
- **IPv6 enabled for most subnets**
  - IPv6 done: 719
  - IPv6 investigate: 24 (wireless etc)
  - IPv6 prohibited: 31
- **~25% of DNS address lookups are for IPv6**
- **~6% of Monash's Internet traffic is IPv6**

Addressing Plans

- **IPv4 addressing plans become quite complicated due to the need to reduce wasted address space by micromanaging the number of used and unused addresses per subnet**
  - Lots of work splitting, merging, renumbering, using secondary address ranges, applying to APNIC for another range ...
- **There are many more IPv6 subnets than IPv4 subnets**
  - Every organisation can have a IPv6 /48 – 64k subnets
  - APNIC One-Click IPv6 gives you a IPv6 /32 – as many subnets as the IPv4 Internet has hosts
  - Each IPv6 subnet can have 2^64 (effectively infinite) hosts, or only 1 – it doesn't matter any more
- **IPv6 addressing plans can be quite regular and sparse**

Monash's main staff IPv4 /16

Each square is a /24

Each dot is a single host, coloured by Department

Hard work when a subnet grows and needs a bigger area

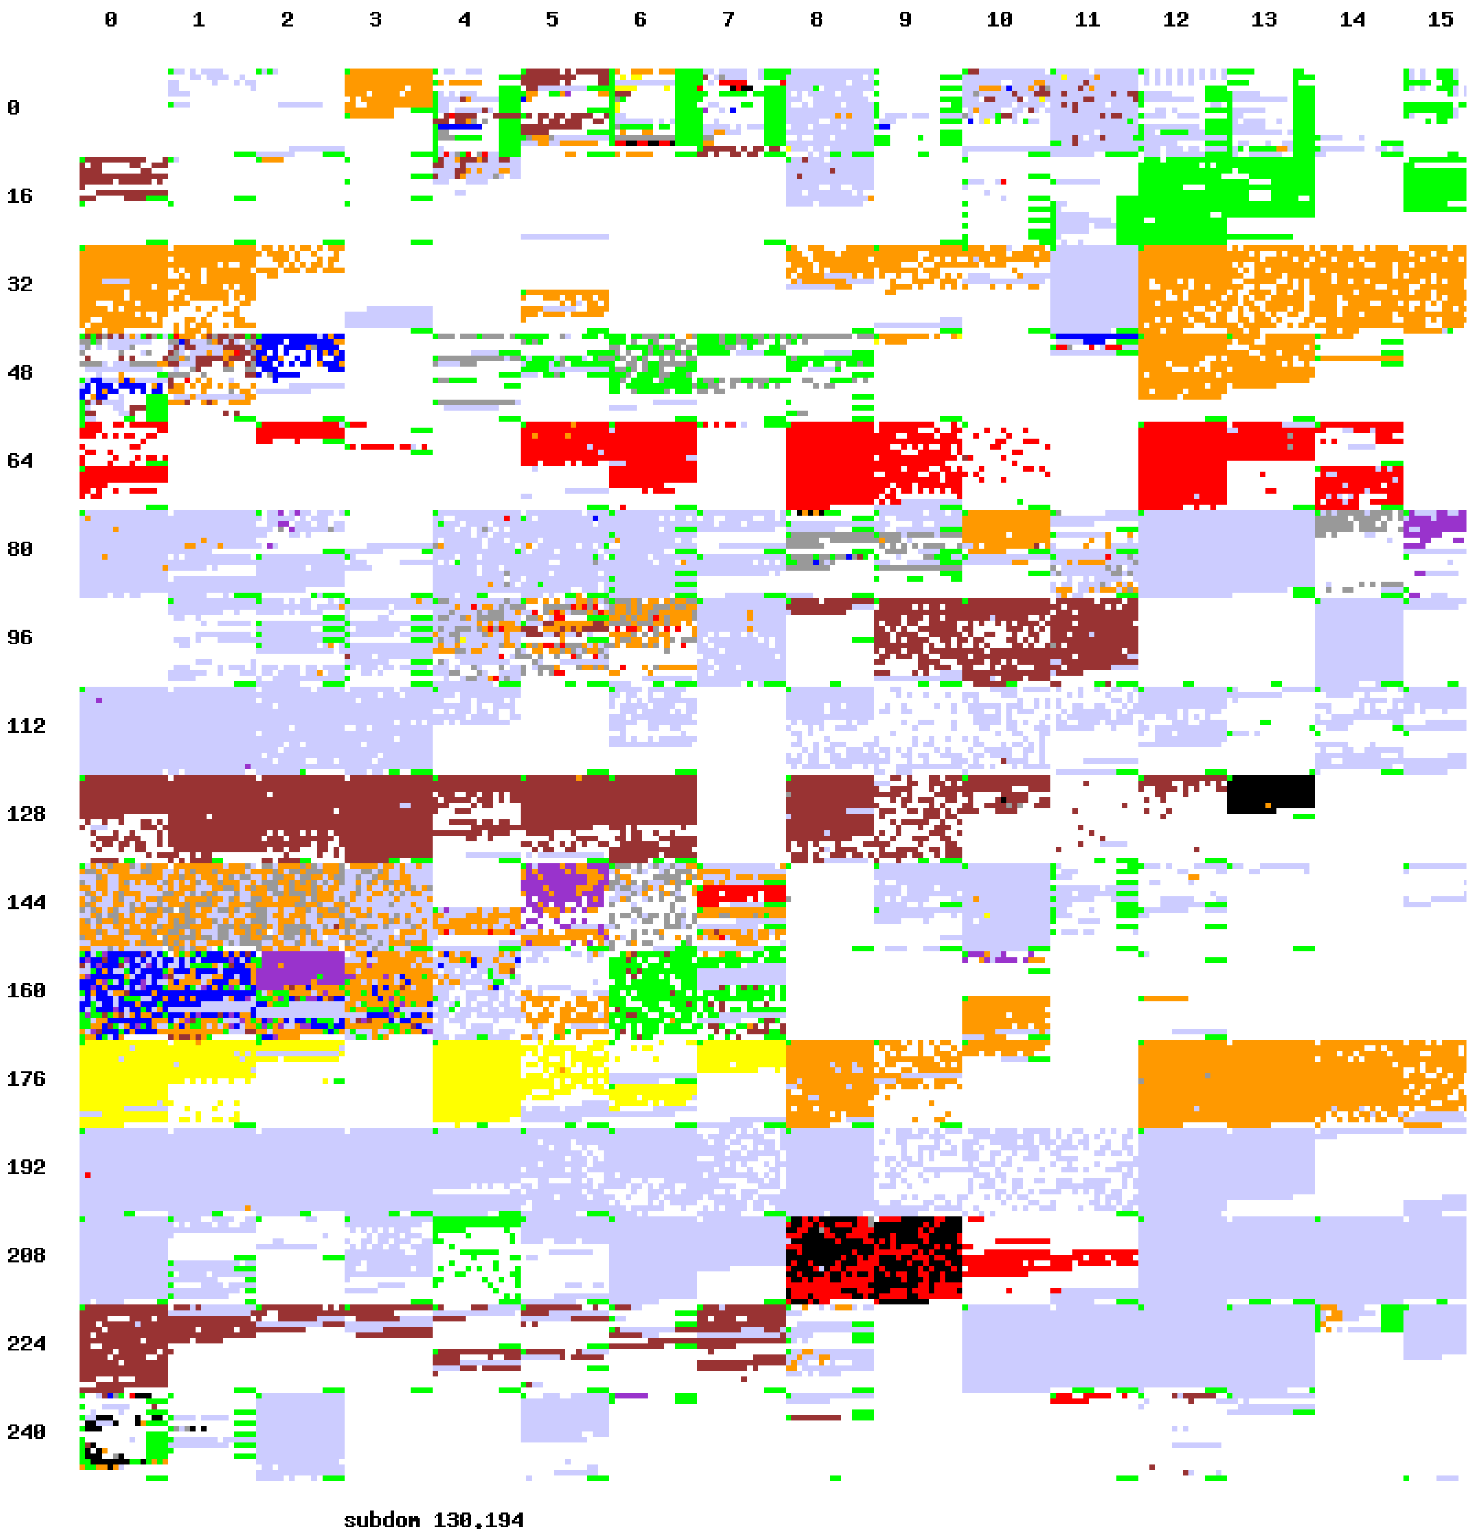

Monash's IPv6 Address Plan

- **Hierarchical addressing plan for ease of summarising ACLs**
  - 4 bits: server / research / staff / student
  - 4, 8 or 12 bits: Organisation unit
  - 8, 4 or 0 bits: Location
- **Backbone links have own /64**
- **Router loopbacks have own /64**
- **Use Unique Local IPv6 Unicast Addresses (FC00::/7), not Site-Local Addresses (FEC0::/10). See RFC 4139, RFC 3879**
- **There are other alternatives, see RFC 5375, RFC 3531, RFC 4057 etc**

Monash's public IPv6 /48

Each square is a /56

Each dot is a /64 subnet which could have 1..2^64 hosts

Multi-location departments get a /58 for each of servers, research, staff, and students

Lots of room for extending the address plan!

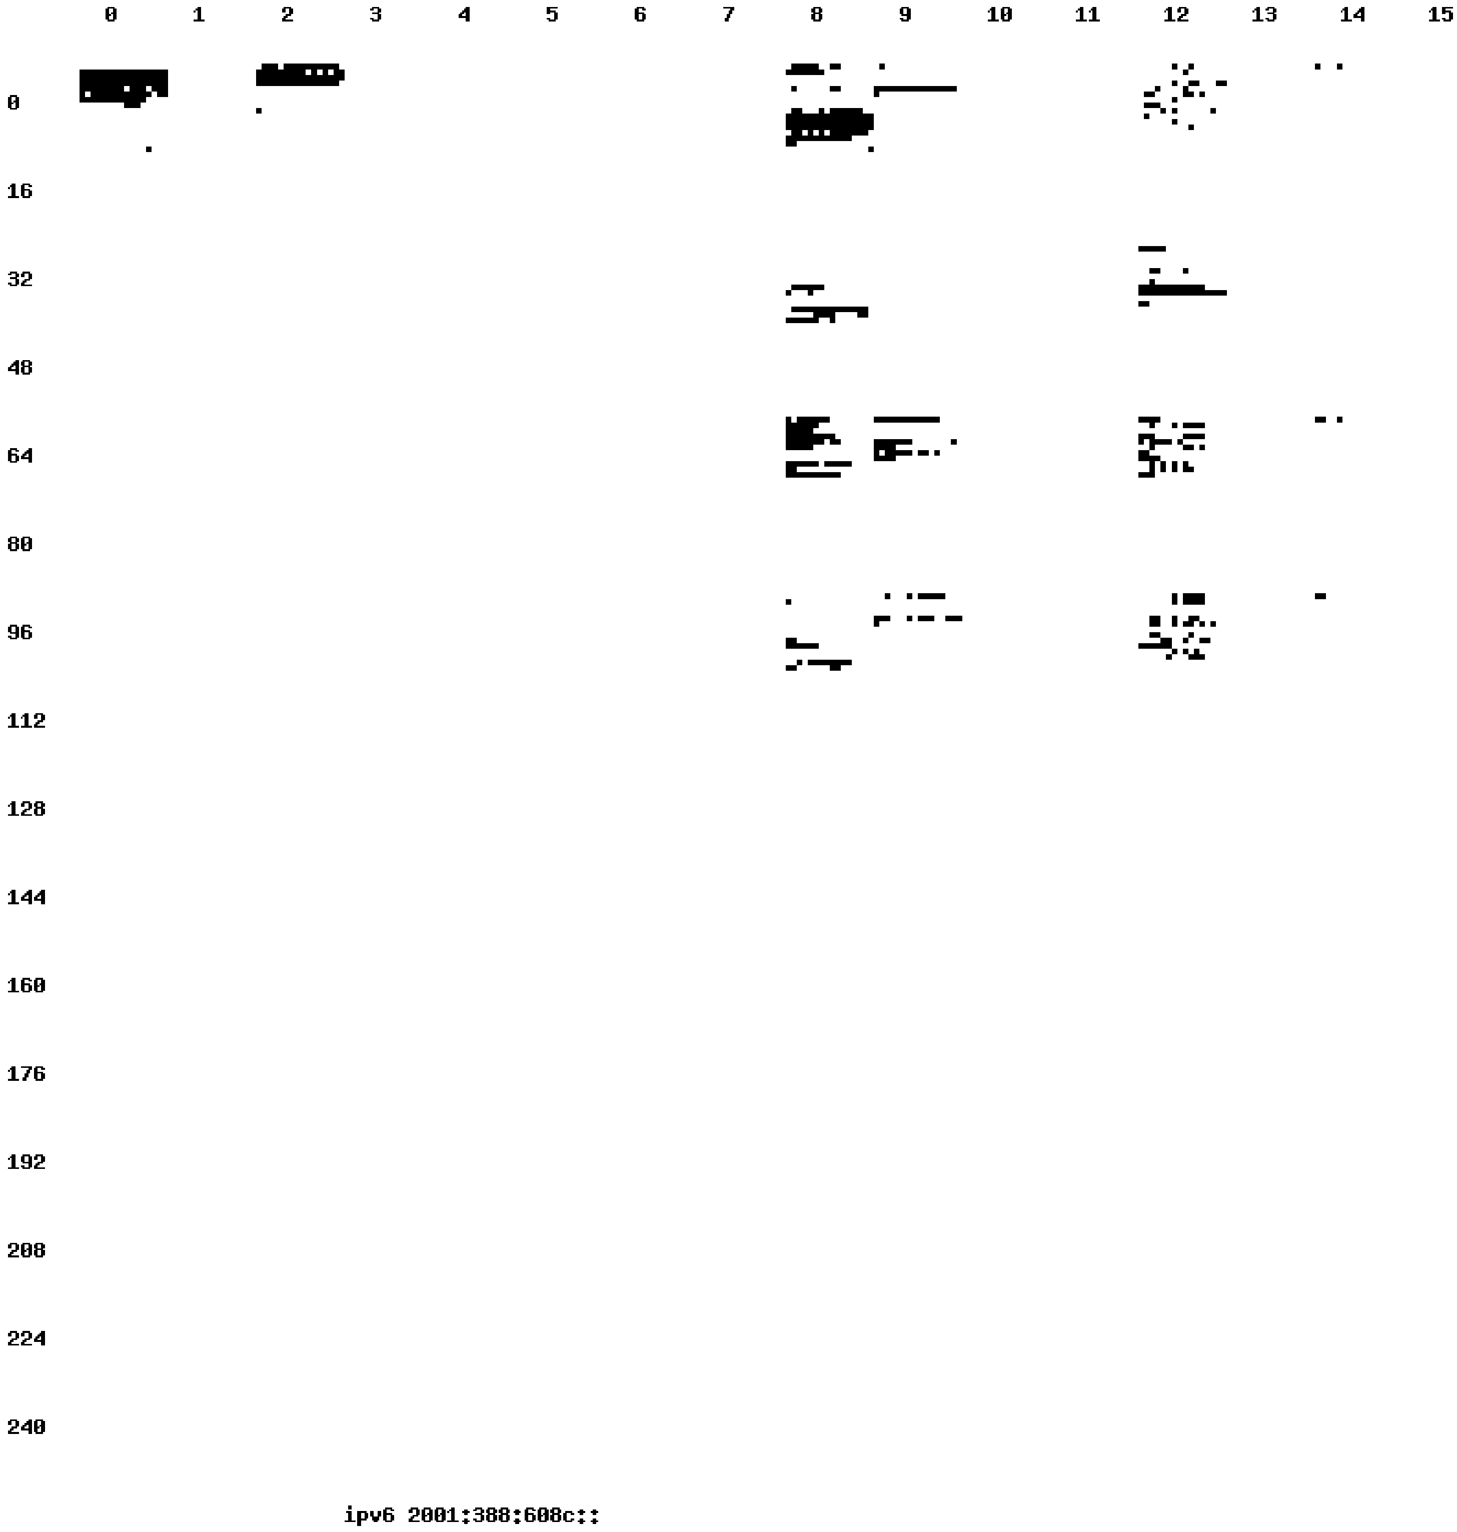

Internet Access Scheme

1. Normal subnets

Have Public IPv4 addresses

Have Public IPv6 addresses

Have direct access to/from Internet

- But limited by policy using ACLs and SCEs

2. Internal subnets

Have Private RFC-1918 IPv4 addresses

Have Private RFC-4193 IPv6 Unique Local Addresses

No Access to/from Internet

- No NAT or Proxies …

We don't mix Public and Private addresses or give different Internet access capabilities for IPv4 v. IPv6

Possible Future Internet Access Scheme

Type 1 and 2 subnets from previous slide, plus

1b. New client subnets

No IPv4 addresses

Have Public IPv6 addresses

Have direct access to/from IPv6 Internet

- But limited by policy using ACLs and SCEs

Use DNS64 / NAT64 to give access to IPv4 Internet

- ? Access control ?

IPv4 v. IPv6 address Usage 2009/2010

- IPv4 ARP table

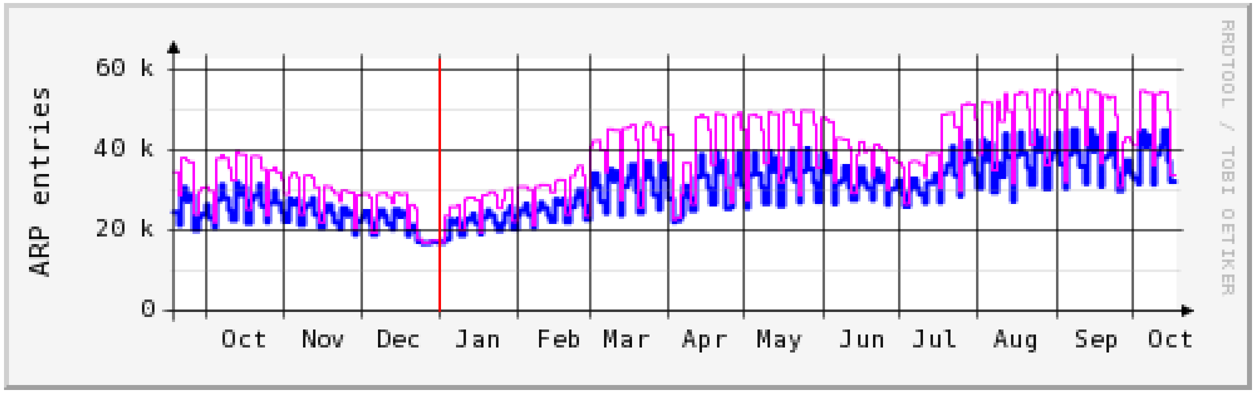

- IPv6 Neighbour table

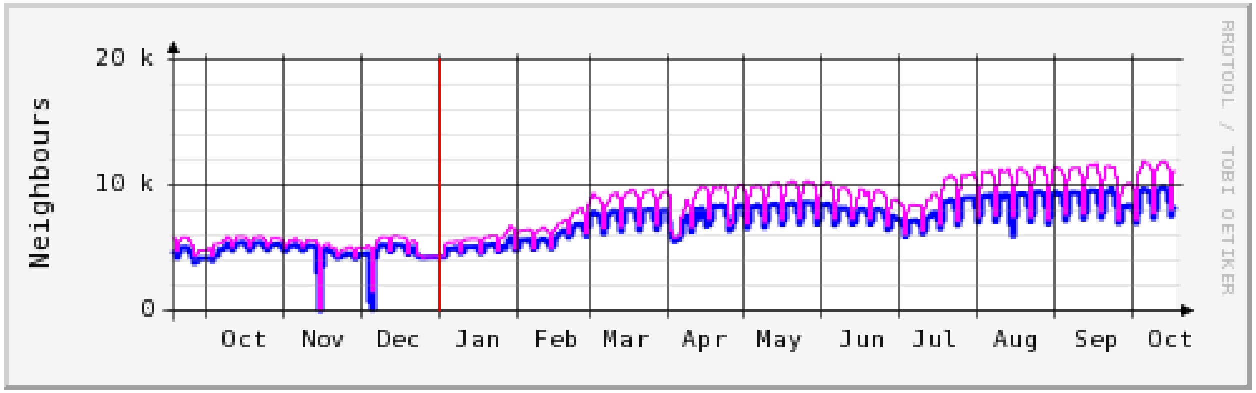

IPv6 is growing and ~20% of IPv4

Slow Progress with End Systems

- **Monash Windows XP desktop SOE now has IPv6 enabled**
- **New Monash Windows 7 desktop SOE (with IPv6 enabled by default) to be released soon**
- **Server owners are still reluctant to actually put the IPv6 address of their servers in the DNS so that clients know to request services over IPv6. Everything works over IPv4, but there _might_ be problems over IPv6**
- **People think "job done" when IPv4 works, don't think to take the extra step to make IPv6 work too**
  - Only think to ask for IPv4 security exemptions
- **Still too early for DHCPv6**

Recommended Commands for Windows

For Windows XP, Start → Run → cmd

```
netsh interface ipv6 install
  (takes a little while)

netsh interface ipv6 set privacy state=disabled
netsh interface ipv6 set teredo type=disabled
netsh interface ipv6 isatap set state disabled
netsh interface ipv6 6to4 set state state=disabled
```

For Windows 7, run “cmd” as Administrator, and extra command

```
netsh interface ipv6 set global randomizeidentifiers=disabled store=persistent
```

Recommended steps for servers

- **Check IPv6 available in O/S**
- **Check IPv6 enabled on network interfaces**
- **Check IPv6 permitted in firewall rules (if any)**
- **Check all applications listen on IPv6 ports**
- **Check applications' access controls (if any)**
- ***Test services using numeric IPv6 address***

No possible disruption to any clients down to this stage

- **Register IPv6 address in DNS using addhost**
  - "IPv6 Subnet: auto"
- ***Test services over IPv6 using hostname***

- **Easy. Mostly things "Just Work"**

Mission-Important IPv6

- **Monash now uses Google Apps for student and staff e-mail, calendar, ...**
- **Monash is part of AARNet's Google over IPv6 membership**
- **When IPv6 (internal or external) doesn't work, students/staff have delays getting to their e-mail**
  - Phones ring
- **But also some resilence**
  - Sometimes IPv6 doesn't work, but IPv4 keeps working
  - Sometimes IPv4 doesn't work, but IPv6 keeps working

Monitoring Network Addresses: IPv4

- At Monash, with IPv4, hosts need to pre-register their MAC address (and optionally an IPv4 address), or authenticate using 802.1x, before getting an IPv4 DHCP lease
- Each host will have only 1 IPv4 address at any one time
- Can track dynamic users using DHCP logs, and RADIUS accounting logs from WISMs or SCEs

Monitoring Network Addresses: IPv6

- Hosts mostly use Stateless Address Autoconfiguration to obtain their IPv6 address (and default gw)
- So, there is no event that flags that a user has joined a subnet
- Users registered on a different subnet, or not registered at all, can get IPv6 addresses, and start using the network

Monitoring Network Addresses: IPv6 (2)

- **Generally, each host will have 2..4 addresses**
  - IPv6 Link-Local address
  - IPv6 stateless autoconfiguration address (RFC 4862)
  - Windows boxes by default will have a IPv6 Temporary Addresses (RFC 4941) or two
  - Routers and servers should also have a static IPv6 address if the address needs to be hard-coded somewhere else
    - Recommend hard-coding IPv6 address for Catalyst 3750 switches since their Ethernet address can change after a reboot.
- **A more-complicated database is required to track, query, display all these addresses**

Need to find the name of host of IPv6 traffic

- **We aren't putting IPv6 forward addresses in the DNS for all clients or servers that speak IPv6**
- **But, we can automatically populate the reverse DNS to make it easier to identify who is sending particular traffic:**
  - Link-local IPv6 address
  - Stateless Autoconfiguration IPv6 address
  - Possibly also IPv6 Temporary Addresses learnt from the Neighbour Discovery tables on the routers
- **Current counts:**
  - Forward DNS: 4026
  - Reverse DNS (Link-Local): 7071
  - Reverse DNS (Global): 8025
  - Reverse DNS (ULA): 1738

Provisioning Systems

- **In general, adding IPv6 is a good opportunity to revisit all your existing network configuration, management, monitoring, procedures and control systems**
  - Start this process with plenty of time, before you need to deploy IPv6 in a last-minute rush
  - IPv6 addresses are long. You do **NOT** want to be typing them by hand, also want to avoid cut-and-paste errors.
  - Make a computer take care of all the drudgery of creating and applying configurations, they are good at menial tasks
- **We needed to extend our systems to allocate IPv6 addresses, create IPv6 router configs and ACLs**
  - Changed from referring to a subnet by their IPv4 address, to referring to it by name
  - ACLs are now named using the subnet name, not IPv4 address

Monash University Victorian Network

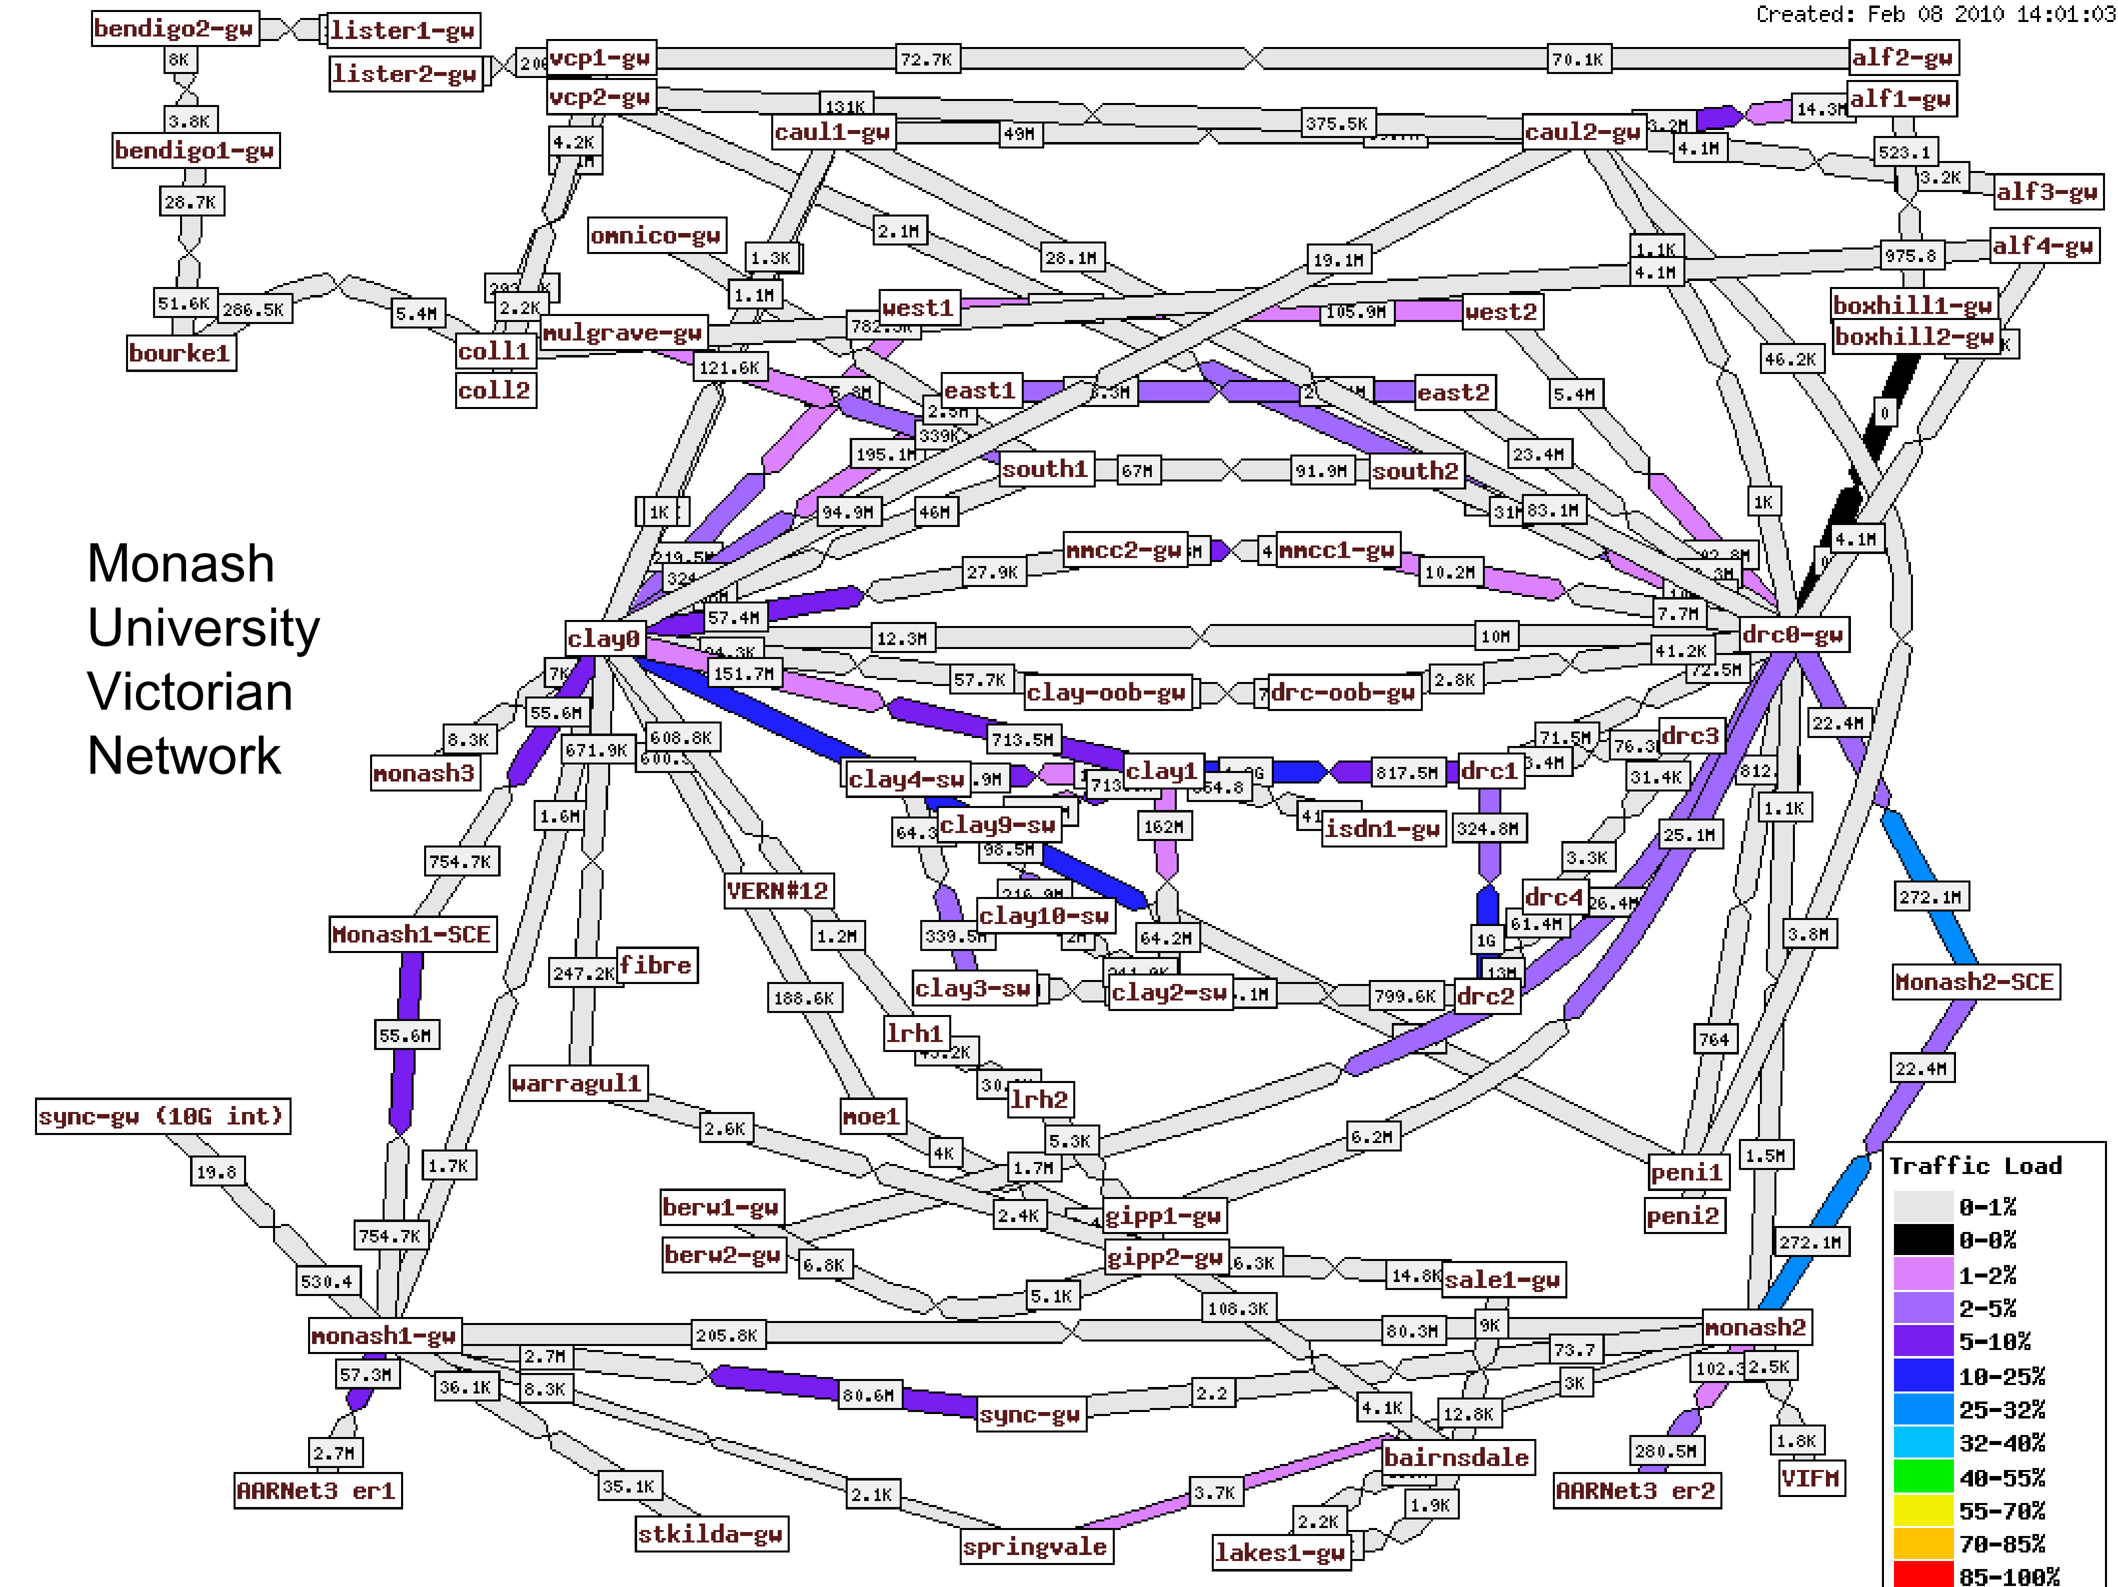

Network Management Tools

- Monash is fortunate enough to have a large enough network, and enough skilled staff, that writing our own network tools is cost-effective
- Don't have to wait for vendors
- Get tools customised to do exactly what we want

NetFlow Traffic Statistics

- **Need NetFlow V9 to collect IPv6 statistics**
- **Securtity Team use “flow-tools” which are V5 only**
  - Added a “flowd2ft” translator front-end
- **Fluke NetFlow Tracker handles V9**
  - But ignores any IPv6 flows
- **Networks Team use NfSen / nfdump (from Sourceforge) to collect and analyse IPv6 traffic data**
- **Other scripts to graph daily counters etc**
- **Note: No NetFlow if traffic not CEF-switched, e.g. IPv6 inbound on Cisco 877 ADSL interfaces**

YouTube content over IPv6 Jan 29 2010

- Large jump in IPv6 data from Google 2001:4860::/32

Was 2 GB/day

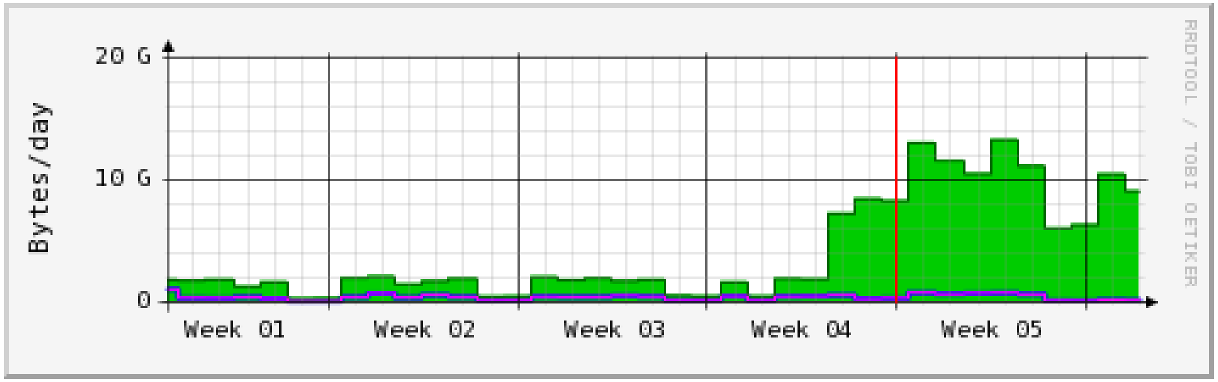

- YouTube data now comes over IPv6 2001:4860:4001::/48

Extra 10 GB/day

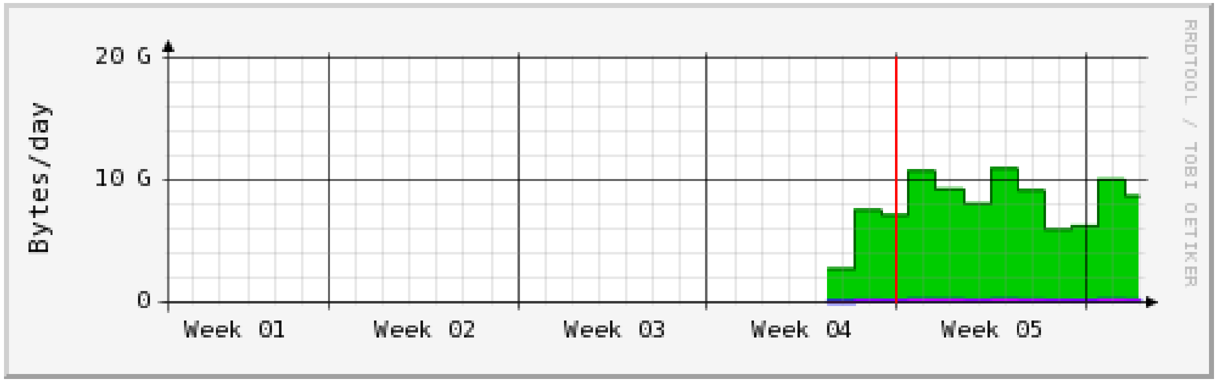

IPv6 Internet traffic Sep/Oct 2010

- IPv6 Internet traffic averages about 170GB per day

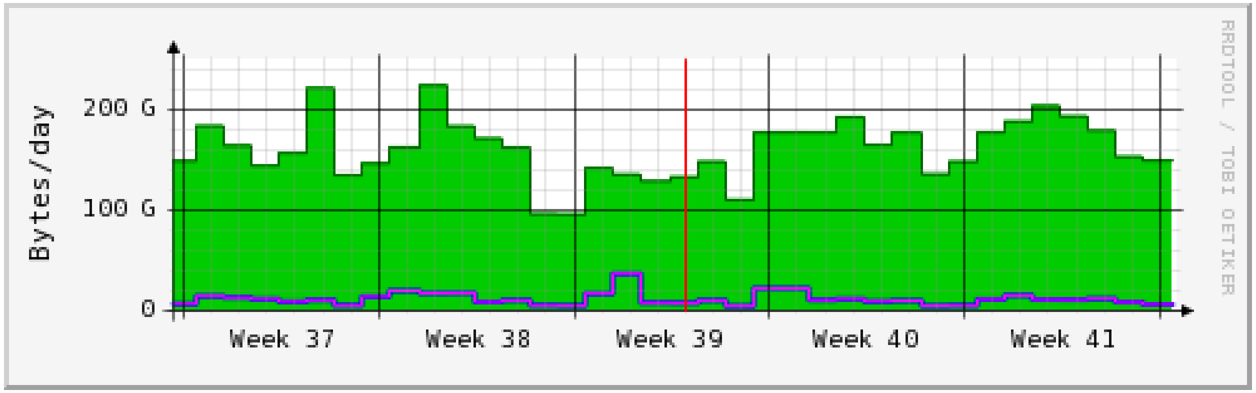

- Number of hosts inside Monash that send IPv6 Internet traffic

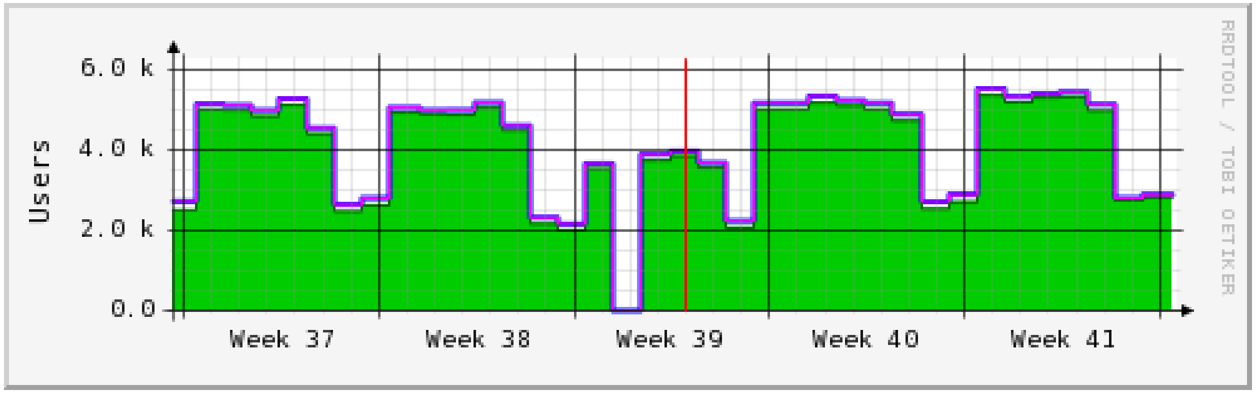

(drop-off due to exams and gap due to disk full)

IPv6 Web server traffic 2009/2010

IPv6 hits/day on central Web farm.

Second or Third most-popular Education Web site that is IPv6 enabled

http://bgp.he.net/ipv6-progress-report.cgi

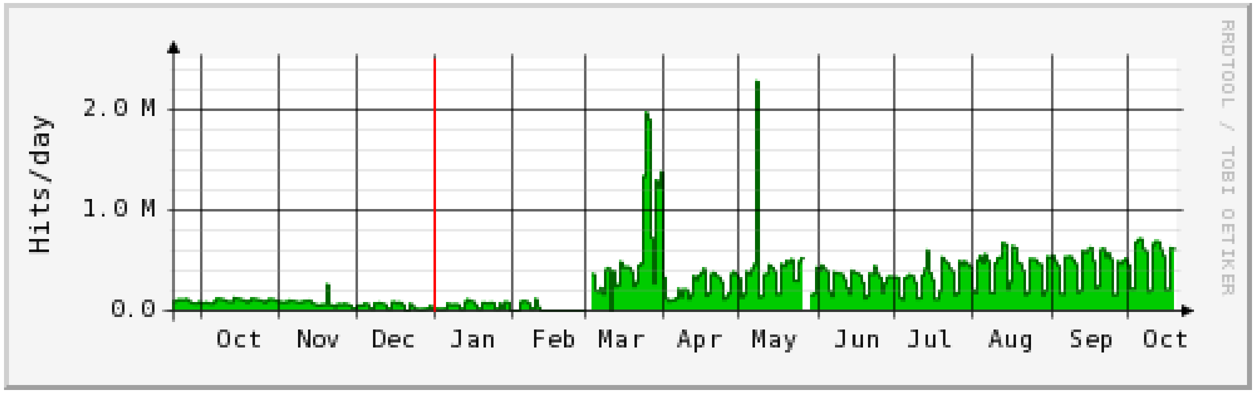

Percentage of hits on {www.}monash.edu{.au} that came over IPv6

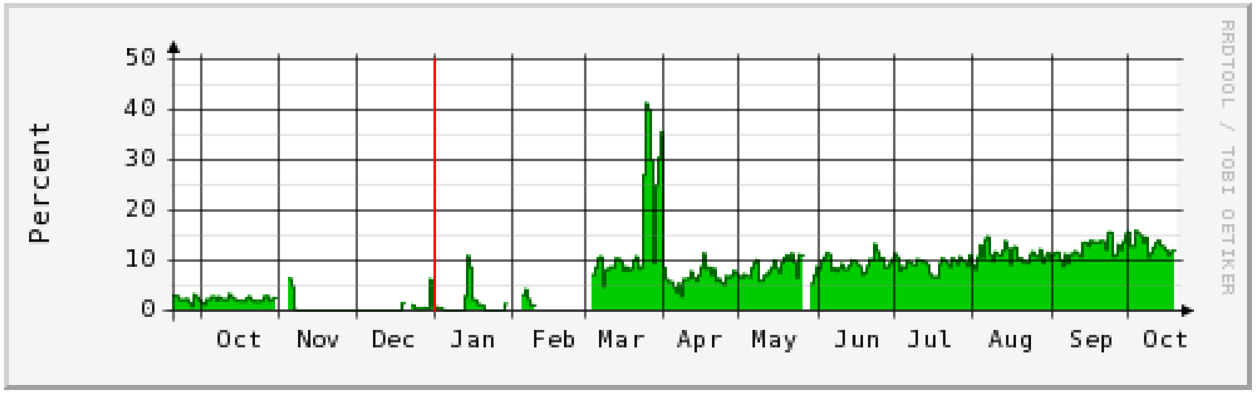

(spike due to web spider.)

IPv4 v. IPv6 traffic at Monash border

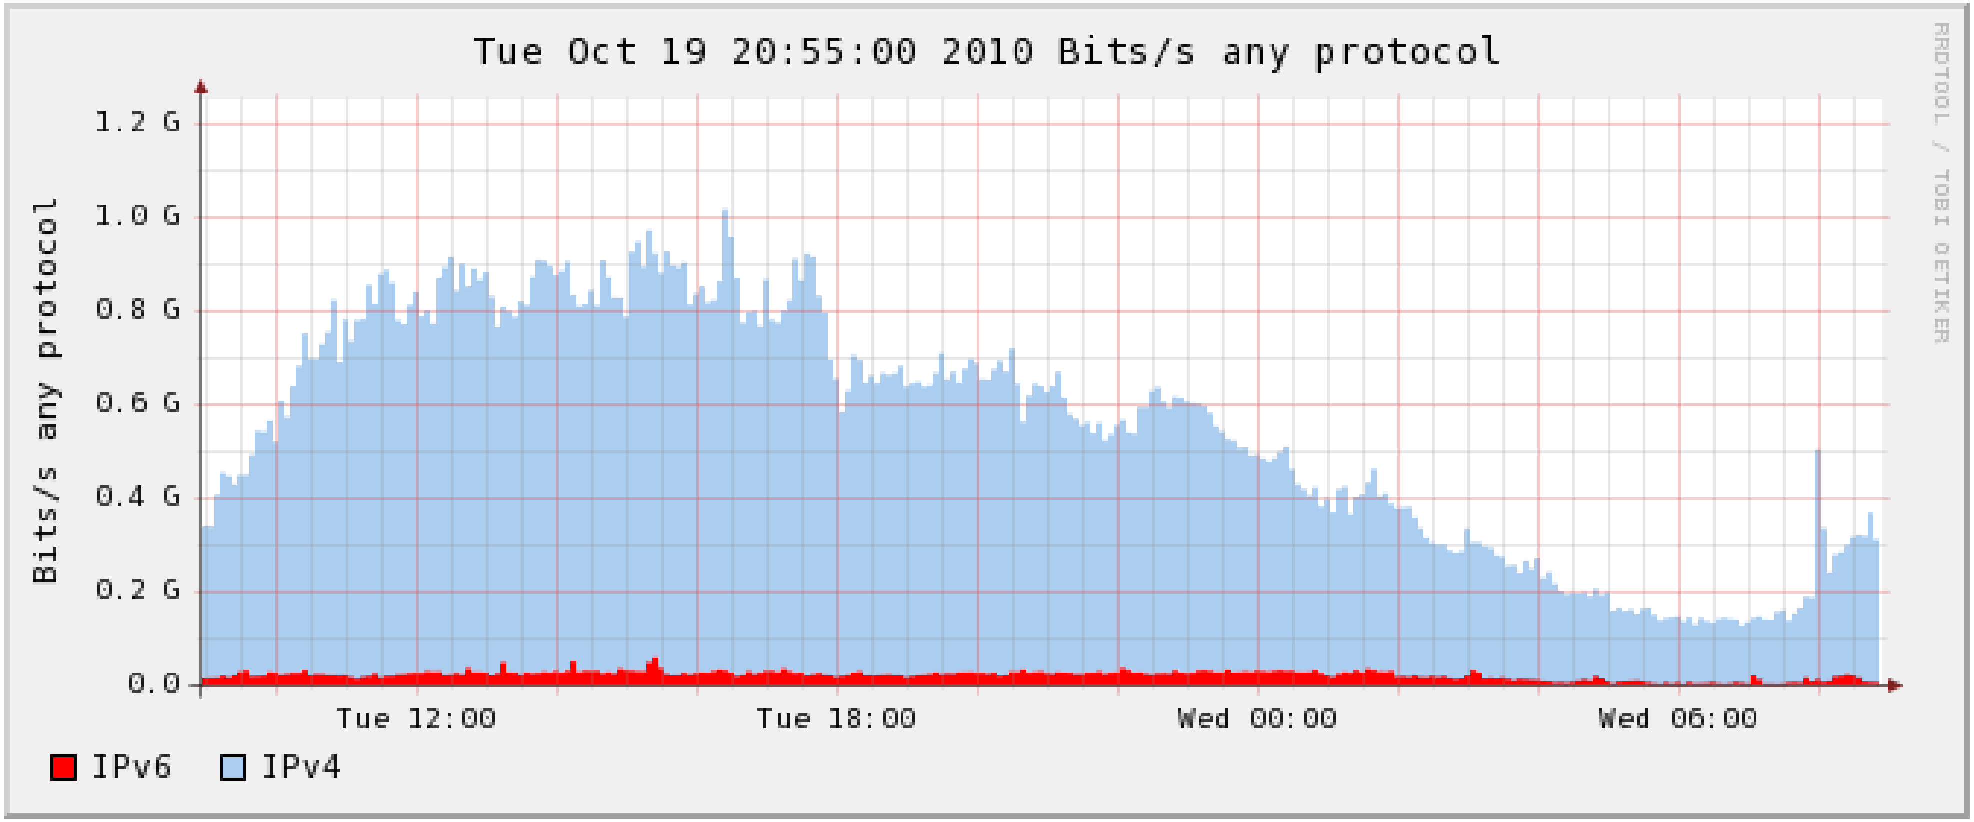

IPv4 average ~500 Mbit/s IPv6 average ~20 Mbit/s

IPv6 Reachability and Delay

- We use SmokePing's Ping6 probe http://oss.oetiker.ch/smokeping/ for both IPv4 and IPv6 delay monitoring
- We use StatSeeker for interface up/down and IPv4 reachability, but not IPv6
- Quick hack script using ping6 to monitor IPv6 reachability

mail.google.com IPv6 now in Sydney SmokePing graph 5 Feb 2010

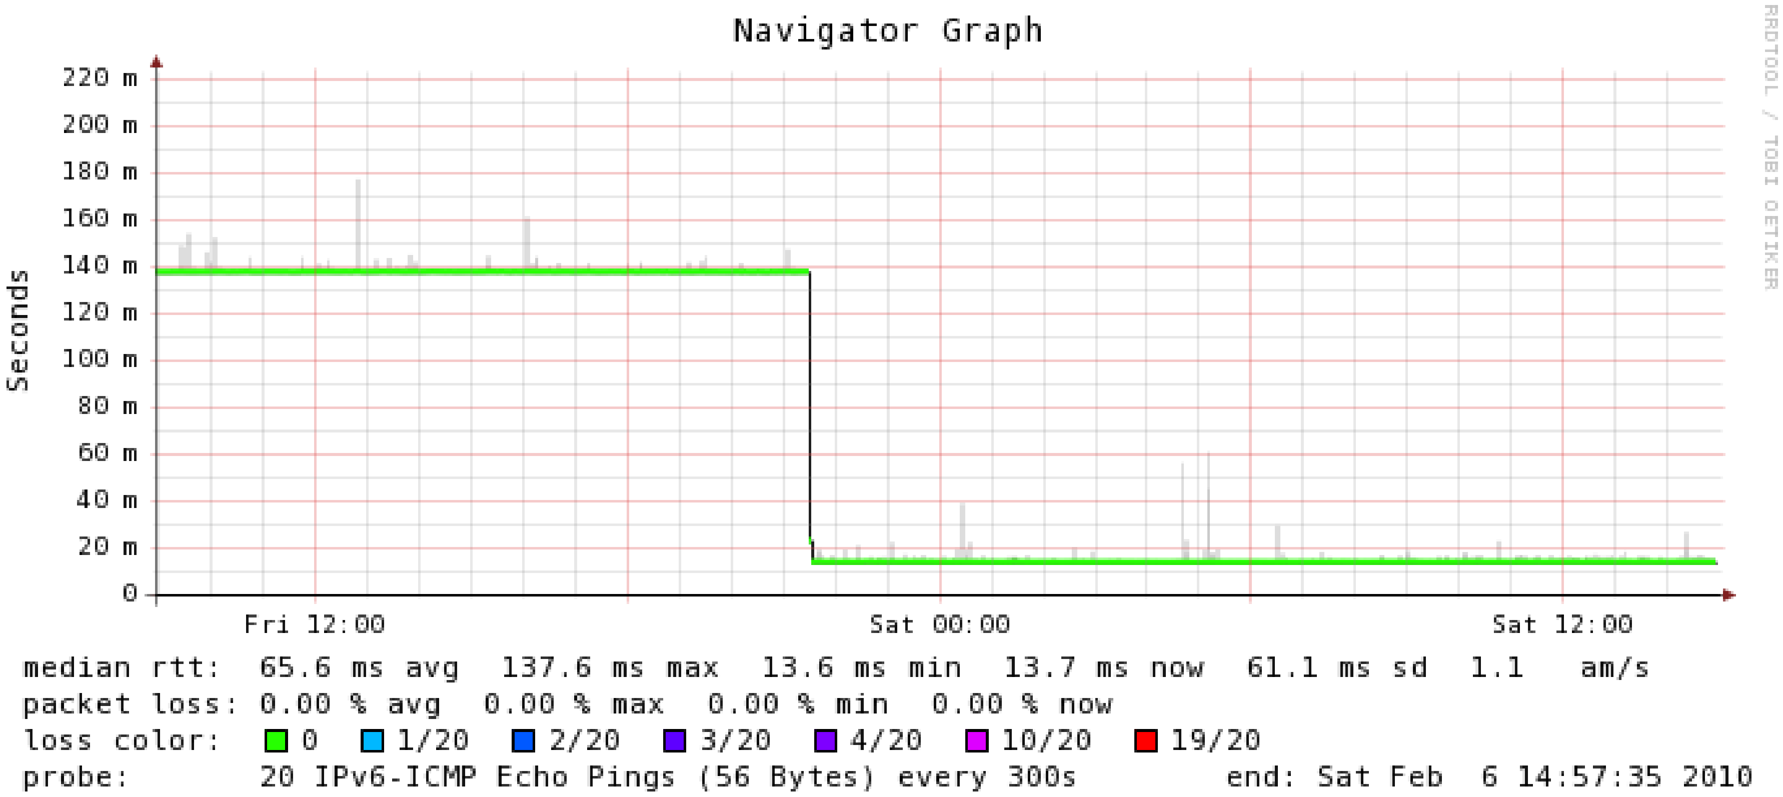

```
median rtt:   65.6 ms avg   137.6 ms max   13.6 ms min   13.7 ms now   61.1 ms sd   1.1    am/s
packet loss:  0.00 % avg   0.00 % max   0.00 % min   0.00 % now
loss color:   🟩 0   🟦 1/20   🟦 2/20   🟪 3/20   🟪 4/20   🟪 10/20   🟥 19/20
probe:        20 IPv6-ICMP Echo Pings (56 Bytes) every 300s        end: Sat Feb  6 14:57:35 2010
```

Reachability

```
Wed Oct 20 09:42:16 EST 2010
1937 things being monitored, 6 things down

fdfd:eb1a:eb14:1400::aa clay-11e-760-rk1-fes1.net.monash.edu
fdfd:eb1a:eb14:1c00::ad peni-1abloomst-house2-rk1-fes1.net.monash.edu
fdfd:eb1a:eb14:1c00::ae peni-1abloomst-house3-rk1-fes1.net.monash.edu
fdfd:eb1a:eb14:1c00::97 peni-1bloomst-unit1-rk1-fes1.net.monash.edu
fdfd:eb1a:eb14:1c00::a peni-1bloomst-unit5-rk1-bds1.net.monash.edu
fdfd:eb1a:eb14:1c00::98 peni-1bloomst-unit9-rk1-fes1.net.monash.edu
```

Vendor Support for IPv6 is often incomplete

- **Cisco CSM doesn't support IPv6**
  - Can route IPv6 (not load-balanced) around CSM
  - Looking at replacement as part of New Datacentre and/or Gen5 Network Projects
- **Cisco WISM doesn't support IPv6 well**
  - Vlan steering of IPv6 traffic works in one direction, not both
  - In short term, maybe special SSIDs for IPv6
- **IPv6 Internet Authentication and control**
  - Upgraded to SCE 8000, software soon
- **Demand "IPv4/IPv6 Feature and Performance Parity"**
  - Every feature (and combination of features) supported in IPv4 needs to be supported in IPv6 at same speed
- **"IPv6 Compliance" is not enough!**

Problem: Increased TCAM usage

- **Increased use of TCAM space since we are managing both IPv4 and IPv6**
- **Catalyst 3750G's**
  - IPv6 features now in IPBASE
  - Enable using “sdm prefer dual-ipv4-and-ipv6 vlan”
    - Allocates about half the TCAM space for IPv6
  - Can't do ACL masks longer that 64 bits
  - Each IPv6 entry takes up twice the space of an IPv4 entry
- **Catalyst 6500's**
  - IPv4 and IPv6 share same TCAM space
  - Configure “mpls ipv6 acl compress address unicast”
    - Ignores bits 39..24 of IPv6 address, which are normally FF:FE or 00:00
    - Drastically reduces amount of TCAM used
- **Avoid individual IPv6 host address ACLs**
  - Can't track privacy addresses fast enough
  - Can't filter /128's on 3750s
  - Ignoring bits 39..24 on 6500s

Problem: Voice Vlan RA Leakage

- **IPv6 configured on routers for most subnets**
  - Router ACLs including ingress anti-spoofing
- **Catalyst 3750 user ports configured for**
  - Data Vlan (untagged traffic)
  - Voice Vlan (802.1Q typically tagged with Vlan 729)
- **If Switch → Cisco Phone → PC**
  - Cisco Phone uses the Voice Vlan
  - Only passes untagged Data Vlan traffic to PC
- **If Switch → PC**
  - PC gets both Data Vlan and Voice Vlan traffic
  - PC sees IPv6 RA (Router Advertisements) on both subnets
  - May choose to use RA from Voice Vlan to auto-configure IPv6 address
  - PC sends all traffic un-tagged
    - Switch passes this traffic back on to Data Vlan
    - Anti-spoof ACL in router drops the traffic as it has a Voice Vlan source address

Solution: Smart port macros

http://www.cisco.com/en/US/docs/switches/lan/catalyst3750/software/
release/12.2_52_se/configuration/guide/swmacro.html

```
macro auto global processing
Note: ALL uplink ports should have 'no macro auto processing'
macro auto execute CISCO_PHONE_EVENT  {
    if [[ $LINKUP -eq YES ]]; then
        conf t
            interface  $INTERFACE
                macro description $TRIGGER
                if [[ $AUTH_ENABLED -eq NO ]]; then
                   description YAYE - there is a phone here!
                   switchport voice vlan 729
                fi
            exit
    ...
```

- **When phone connects to a port, CDP triggers adding the Voice Vlan, link down triggers removing VoIce Vlan**

Problems with Windows Vista +

- **A Windows box (e.g. a new personal laptop with Vista) will configure itself as a IPv6 6to4 gateway router if it has**
  - More than 1 interface
  - A global IPv4 address
  - Internet connection sharing enabled, and
  - Not logged into a Domain
- **This will provide 6to4 addresses and tunnel routing to all hosts on the same LAN**
  - And creates an IPv6 black hole when the user takes the laptop home each night
- **Windows 7 tries even harder to use IPv6**
  - Any problems with local IPv6 connectivity, it will fall back to Teredo, tunnelled through gateway overseas (poor performance)

Problems with Mac OS X 10.5 +

- **Prefers IPv6 over IPv4**
- **But, doesn't obey RFC 3484 !**
  - When a client has a global IPv4 address, and a tunnelled IPv6 address, should prefer IPv4
  - When a client has a global IPv4 address, and a link-local IPv6 address, should prefer IPv4
- **Hence is tricked by**
  - Rogue Windows 6to4 router (previous slide)
  - Or a Cisco router configured as an IPv6 "client"
- **Apple OS X IPv6 code is very old**

```
$ sysctl -a | grep -i kame
net.inet6.ip6.kame_version: 20010528/apple-darwin
```

- **See http://lists.apple.com/archives/Ipv6-dev**
- **iOS 4 for iPhone / iPad getting IPv6**

ACLs to make life easier for Macs

- **On Catalyst 3750 user edge ports**

```
interface GigabitEthernet1/0/1
  ipv6 traffic-filter RA-INPUT in

ipv6 access-list RA-INPUT
  deny icmp any any router-advertisement
  deny icmp any any router-renumbering
  permit ipv6 any any
```

- **On Cisco router Vlans configured as “client”**

```
interface Vlan767
  ipv6 traffic-filter BLOCK-RS in

ipv6 access-list BLOCK-RS
  deny icmp any any router-solicitation
  permit ipv6 any any
```

Rogue Router Suppression

- **Rogue routers eliminated**
  - Note: Actually just hidden because these PACL hits can't be logged
- **Really want “RA Guard” (cf. DHCP Trust)**

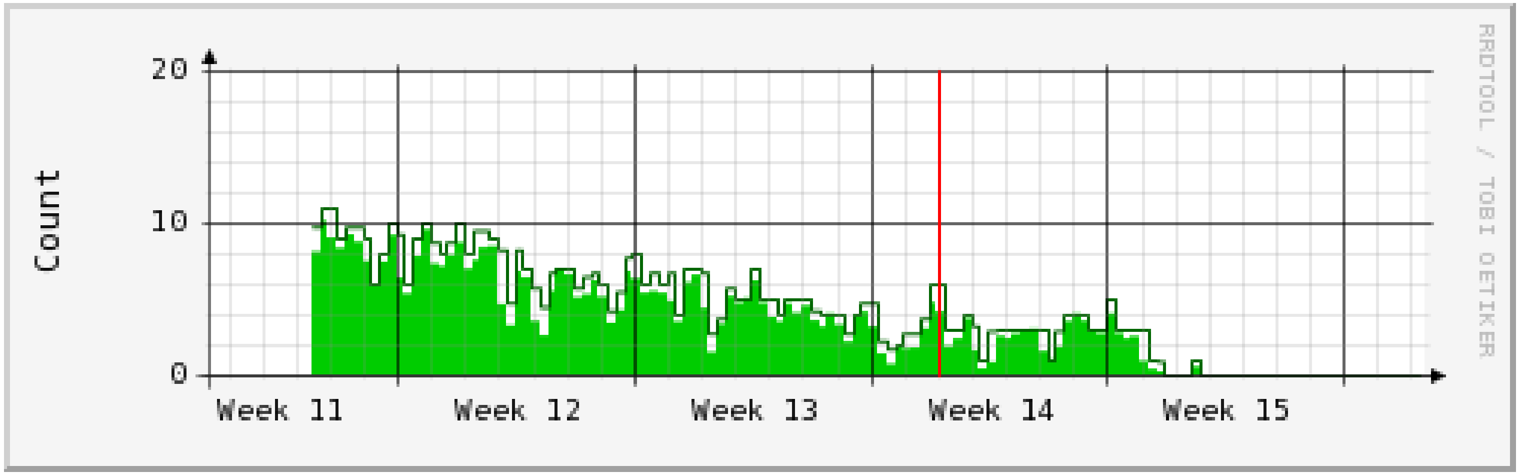

Problem: Windows XP slow YouTube over IPv6

- **By default, for 1-100 Mbit/s transmission, Windows XP sets TCP window size to 17 KB**
  - http://msdn.microsoft.com/en-us/library/ms819736.aspx
- **The YouTube IPv6 server is 185 msec away, so 17 KB window gets <= 85 KB/sec throughput**
  - But Lecturer wanted to play a clip that is 93 KB/sec
- **To solve the bandwidth-delay problem, TCP Window Scaling was defined in RFC 1323 in 1992**
- **But, on Windows XP, enabling window scaling effects IPv4, but not IPv6**
- **Solution: Google plan to enable IPv6 on YouTube servers in Australia in next few months**

What should YOU be doing now?

- **If you haven't started deploying IPv6, start ASAP**
- **Create some sort of IPv6 address plan**
- **Enable IPv6 in your Internet gateway**
- **Get first test server on IPv6**
  - Requires some IPv6 routing and DNS
- **Then, in parallel**
  - Enable IPv6 on Project / Operations / Helpdesk PCs
  - Enable IPv6 on Internet-facing servers
  - Run user awareness and IPv6 training
  - Enable IPv6 on user subnets

“Do What You Can Do First”

- **You don't have time to do a full IPv6 readiness survey, wait for all vendors to add IPv6 to their products, create a complete and comprehensive plan on how to upgrade everything etc etc**

- **Don't worry about**
  - The 5% of systems use an IPv4 thick client to talk to some IPv4 server
  - The 10% of systems that can't have IPv6 enabled until the next upgrade cycle
- **Ignore them until later**
- **There is still plenty to do**
- **Score a few easy goals. Get the easy 50-75% done**
- **Make sure all new / refreshed systems have IPv6 enabled**
  - Much less disruptive than doing it after they go into production

- **Don't aim for perfection first time around. You probably will change address plan, router settings, host settings etc as you gain experience**

Don't be Afraid of Problems

- We have had 20+ years to get all the wrinkles out of our IPv4 networks
- We have 2 years to get IPv6 networking right

- You are going to find things that don't work right. Things that you can't find in a test lab, or by reading the manuals
- You will need to find problems and fix them, before you can uncover the next problem

- It's better to stumble in 2010 (and inconvenience a small percentage of users)
- It won't be OK to stumble in 2012

Risks of not adopting IPv6 now?

- A rushed IPv6 deployment later will be of lower quality and more disruptive
- Extra effort will be required to rework systems already in production
- Risk of forklift upgrades for incompatible systems
- Miss IPv6 opportunities

- Don't hit the IPv4 address exhaustion wall totally unprepared about IPv6 !!

What about NAT?

- **Monash has avoided the use of NAT-P**
  - So we can do traffic accountability and accounting
  - Some current and potential future applications don't work (well) through NAT
  - Hard to do high-speed and resilient NAT
- **Monash has enough Public IPv4s for a couple of years (we hope)**

- **In the long run, we will have Public IPv6-only client subnets**
- **They will still need to access IPv4-only servers**
  - They will need NAT64 / DNS64 until **all** of the Internet supports IPv6

- **Adding any more IPv4 NAT-P now is moving in the wrong direction towards a future IPv6 Internet**

Summary

- **IPv4 addresses are running out**
- **IPv6 is already here**
- **Monash has a good IPv6 infrastructure**
  - Now time to enable for servers and users
- **Vendor support not complete**
  - Do what you can now where you can
- **Start roll-out now to avoid rushing**
  - IPv6 mostly easy to enable and configure
  - Many small simple steps required

Thanks

- **Thank you for listening**
- **Thanks also to**
  - The Network Operations Manager for approving changes to the Monash network
  - Monash Network Engineers for helping with the roll-out and solving problems
  - Support from Management
  - Tolerance from users for occasional problems

Questions?

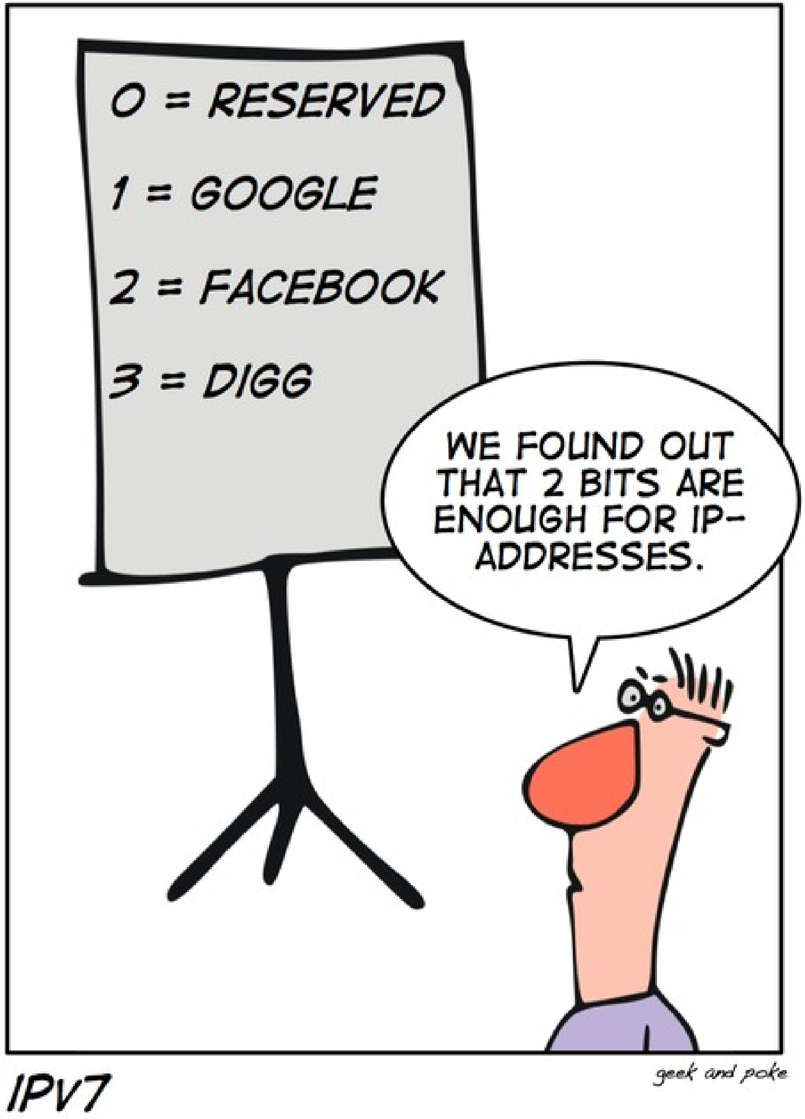
0 = RESERVED
1 = GOOGLE
2 = FACEBOOK
3 = DIGG
WE FOUND OUT THAT 2 BITS ARE ENOUGH FOR IP-ADDRESSES.
IPV7
geek and poke

What can I do? (personally)

- **Get your feet wet with IPv6**
- **Enable IPv6 on your Windows XP desktop and laptop**
  - You shouldn't notice any differences
  - We don't *want* there to be any user visible differences.
  - If there are any problems, report them, so that they can be fixed
- **For bonus points: Set up IPv6 tunnel or native over your home ISP service**

Upgrading the Infrastructure

- **IPv6 is just another infrastructure upgrade**
  - There is no IPv6 “killer app”, except for more addresses to allow the Internet to keep growing.
- **Monash has done infrastructure upgrades before**
  - Terminals → Ethernet → Thinwire → UTP
  - Point-to-point → bus → hub → router → switch
  - Copper wires → multimode fibre → singlemode fibre → CWDM → DWDM
  - DECnet, AppleTalk, Banyan, IPX/SPX → IPv4

Start the IPv6 Roll-out Now

- **Enabling and configuring IPv6 is easy!!**
- **Easier to do roll-out slowly and incrementally**
- **Minimise extra cost and effort by taking advantage of natural replacement cycles**
  - Configure IPv6 during build and test phase
  - Less Change Management !!!
- **Give support staff time to gain knowledge and experience – learn while doing**