World IPv6 Day: Enterprise Sites (and some IETF stuff)

UKNOF 19, Leeds, April 2011

Tim Chown

Electronics and Computer Science (ECS)

University of Southampton

tjc@ecs.soton.ac.uk

Participation?

- We already have dual-stacked internal clients and networks as well as public-facing services
  - Not taking part in Google IPv6 whitelist experiment
- Our interest in World IPv6 Day?
  - Performance under higher IPv6 traffic loads
    - Biggest hit likely to be YouTube over IPv6
  - Measurement of client/application behaviour
    - How many inbound user connections are we losing, and why?
    - What issues do our users have connecting to IPv6 sites?
- Decided to document the broader issues:
  - See draft-chown-v6ops-ipv6-call-to-arms-02

Internet Draft aims

- Raise awareness of the June 8[th] World IPv6 Day
- Seek to capture:
  - Common causes of connectivity and performance issues, focusing on those that an end-site can influence, with suggested actions/measures
  - Methods to measure and monitor IPv6 traffic, to allow analysis of the traffic behaviour
- Produce a final version well in advance of IPv6 Day
  - Hopefully the advice may have longer-term benefits
  - Measurement tools could be left in place
  - Not expected that sites will configure *client* IPv6 access just for the day – longer-term deployment preferable

Connectivity issues?

- Currently cited (in no particular order):
  - Unmanaged tunnels (6to4 relays, proto41 filters,…)
    - CPEs with 6to4 on by default
  - Tunnel broker first-hop
  - Connection timeouts (failover to/from IPv4/IPv6)
  - PMTU discovery (ICMPv6 filtering)
  - Rogue router advertisements (inc. IPv4-only networks)
  - Tunnel performance (under higher load)
  - AAAA record advertised but service not enabled
  - Filtering – impact of drop vs unreachable

More on connection timeouts…

- What happens if IPv6 network connections fail and browsers need to fall back to IPv4?
  - e.g. due to a local rogue RA, or broken 6to4
- Some good stats gathered by Teemu Savolainen (Nokia)
  - Presented at Prague IETF, March 2011
  - http://www.ietf.org/proceedings/80/slides/v6ops-12.pdf
- Tests whether an unreachable indication helps
  - A key point to note is Windows client performance
    - 21 second delay, whichever browser used, whether unreachable received or not
  - Linux does better if given an unreachable

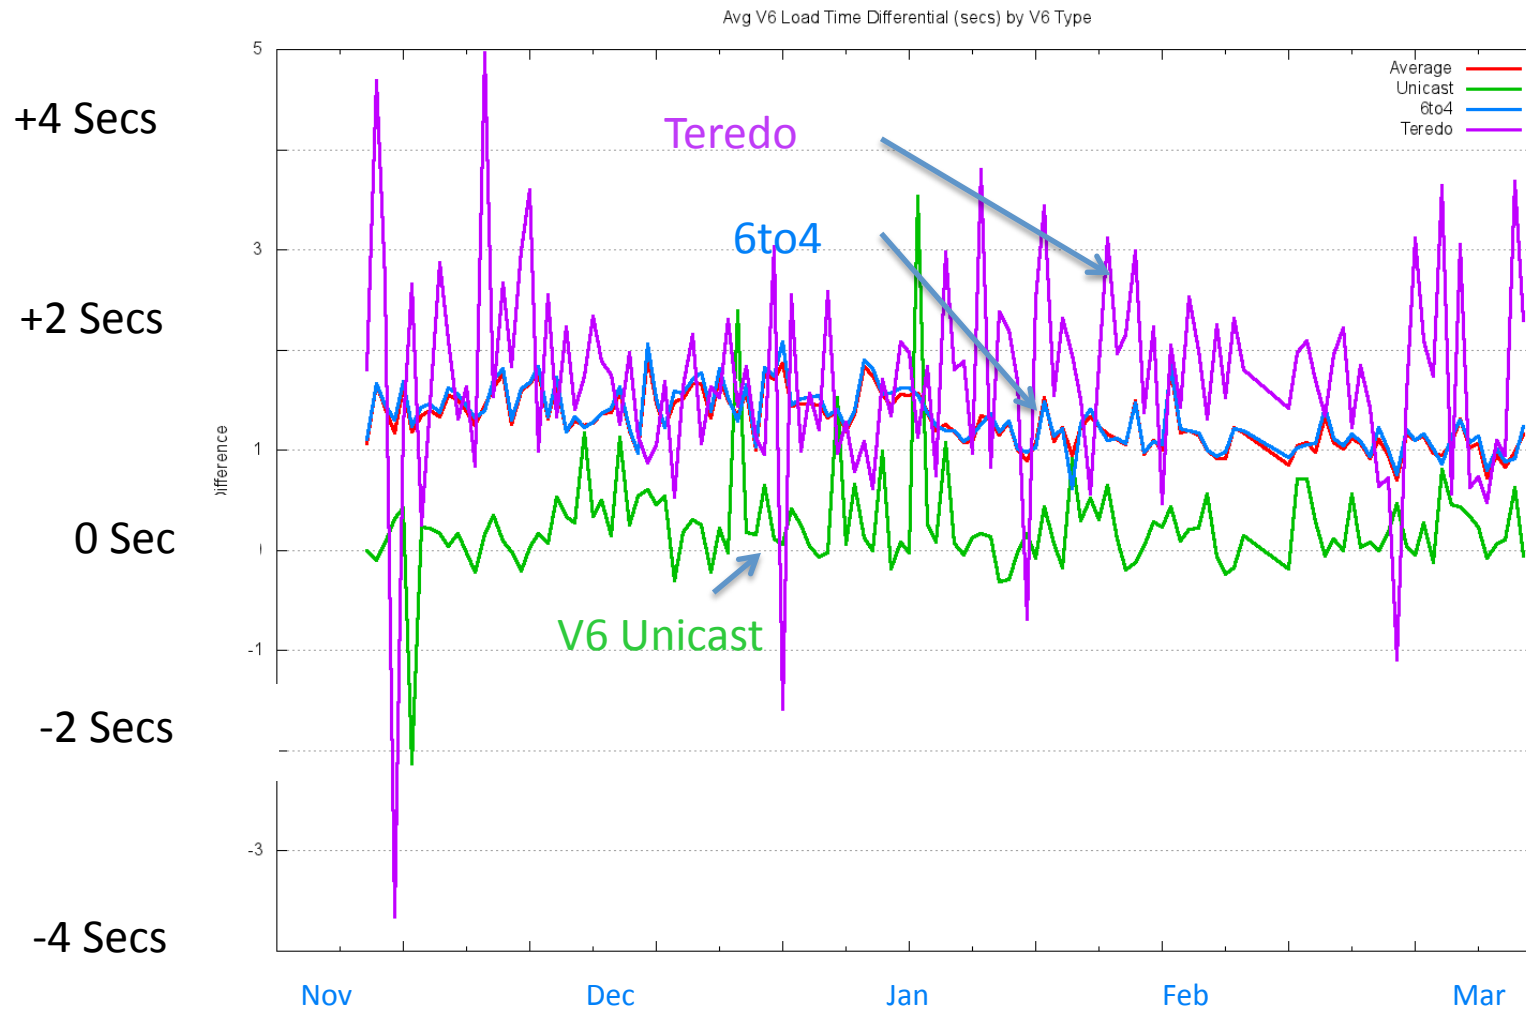| Device	DNS query sending style	IPv6 broken, time until fallback to IPv4			Comments
	Dual-stack destination	Black hole	No route	Address unreachable	
Symbian^3 on Nokia N8 (11.012)	**A** first and used if possible. **AAAA** if no IPv4.	N/A	N/A	N/A	Symbian^3 prefers IPv4 hence tested fallback scenarios are N/A. The DNS query order is a configuration parameter.
Windows 7 Starter Edition on HP *IE 8.0.7600 & Google Chrome 8.0.552.224 & Safari 5.0.2*	**A** and after reply **AAAA**. Uses IPv6 if both available.	**~21**s	**~21**s (after 3 SYN & ICMPv6 errors)	**~21**s (after 3 SYN & ICMPv6 errors)	Same initial delay with those browsers. **The 21 seconds is TCP timeout after 3rd SYN failed.**
iOS4 4.2.1 on Apple iPhone4 *Safari*	**A** first and **AAAA** immediately after. Uses IPv6 if both available.	**No fallback**	**~4**s (After 5 SYN & ICMPv6)	**~4**s (After 5 SYN & ICMPv6)	Lucky observation: waits ~350 ms for AAAA to arrive after A is received before going for IPv4
Apple OS/X 10.6.6 on iMac *Safari 5.0.3 Firefox 3.6.13*	**A** first and **AAAA** immediately after. Uses IPv6 if both available.	**~75**s	**~4**s (After 5 SYN & ICMPv6)	**~4**s (After 5 SYN & ICMPv6) **Firefox:** no fallback at all!	Special note that Firefox did not fallback on address unreachable error.
Android 2.3.1 on Samsung Nexus S *Native browser*	**AAAA** and after reply **A**. Uses IPv6 if both available.	**~21**s	**~0**s (acts on first ICMPv6)	**~0**s (acts on first ICMPv6)	**The 21 seconds is TCP timeout after 3rd SYN failed.**
Maemo5 IPv6 enabled version on Nokia N900 *Firefox & native*	**AAAA** and after reply **A**. Uses IPv6 if both available.	**~189**s	**~0**s (acts on first ICMPv6)	**~0**s (acts on first ICMPv6)	189s is after 6th SYN failed. Kernel: 2.6.28-based
Ubuntu 10.04 / 10.10 on "PC" *Firefox 3.6.13*	**AAAA** and after reply **A**. Uses IPv6 if both available.	**~21**s	**~0**s (acts on first ICMPv6)	**~0**s (acts on first ICMPv6)	Note: immediate fallback to IPv4 happens also during complex page load (i.e. minimizes damage when IPv6 is always preferred) Kernel (10.04): 2.6.32-27, (10.10): 2.6.35-24

Happy Eyeballs

- One way to help clients is to use a 'happy eyeballs' approach
  - Try both IPv4 and IPv6 at the same time
  - Use the connection that works
    - Can favour IPv6 initially by giving it a slight head start
    - Tuned using a 'p' value, if negative, favour IPv4
  - See draft-ietf-v6ops-happy-eyeballs-01
  - Some concerns about extra connection load
    - Would typically drop the 'losing' connection
  - Some interest in using happy eyeballs for multiple interface (mif) scenarios
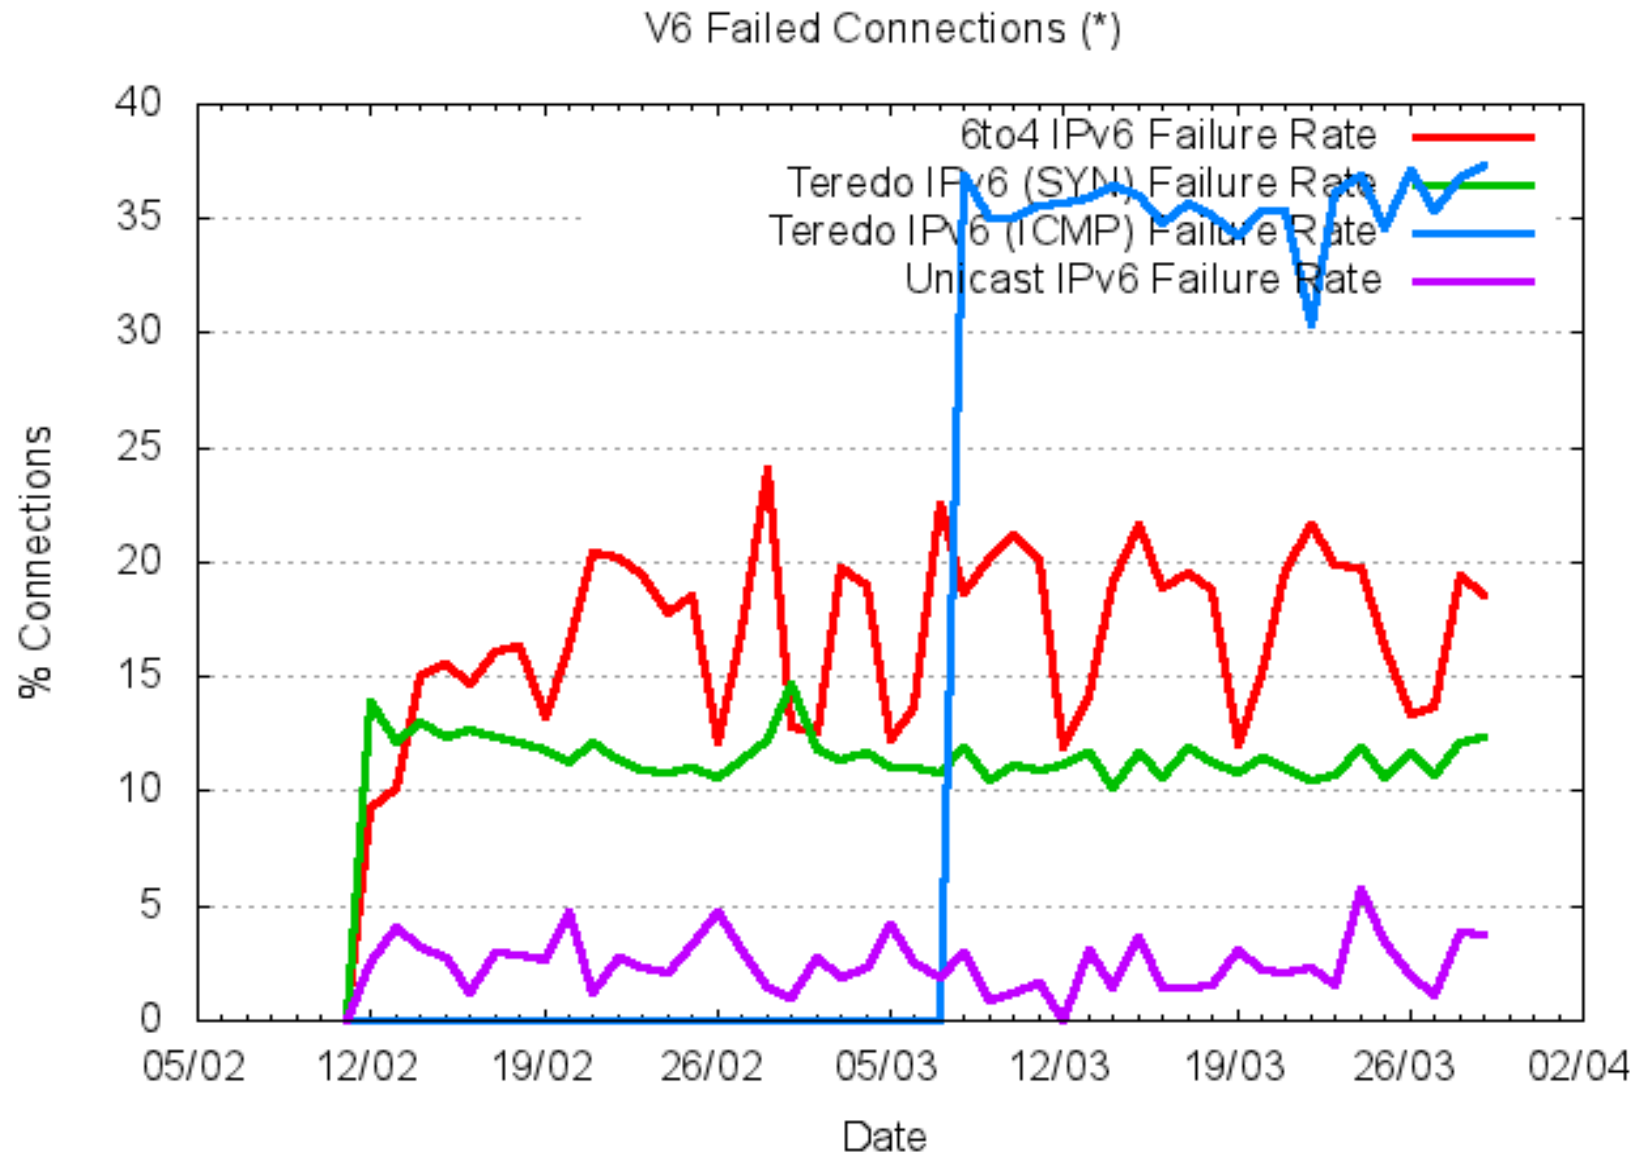
More on tunnel performance

- How well do 6to4 and Teredo perform?
  - Do we really want users using these methods?
- Stats from measurements were presented by Geoff Huston at the Prague IETF, March 2011
  - Shows increased latencies for 6to4/Teredo
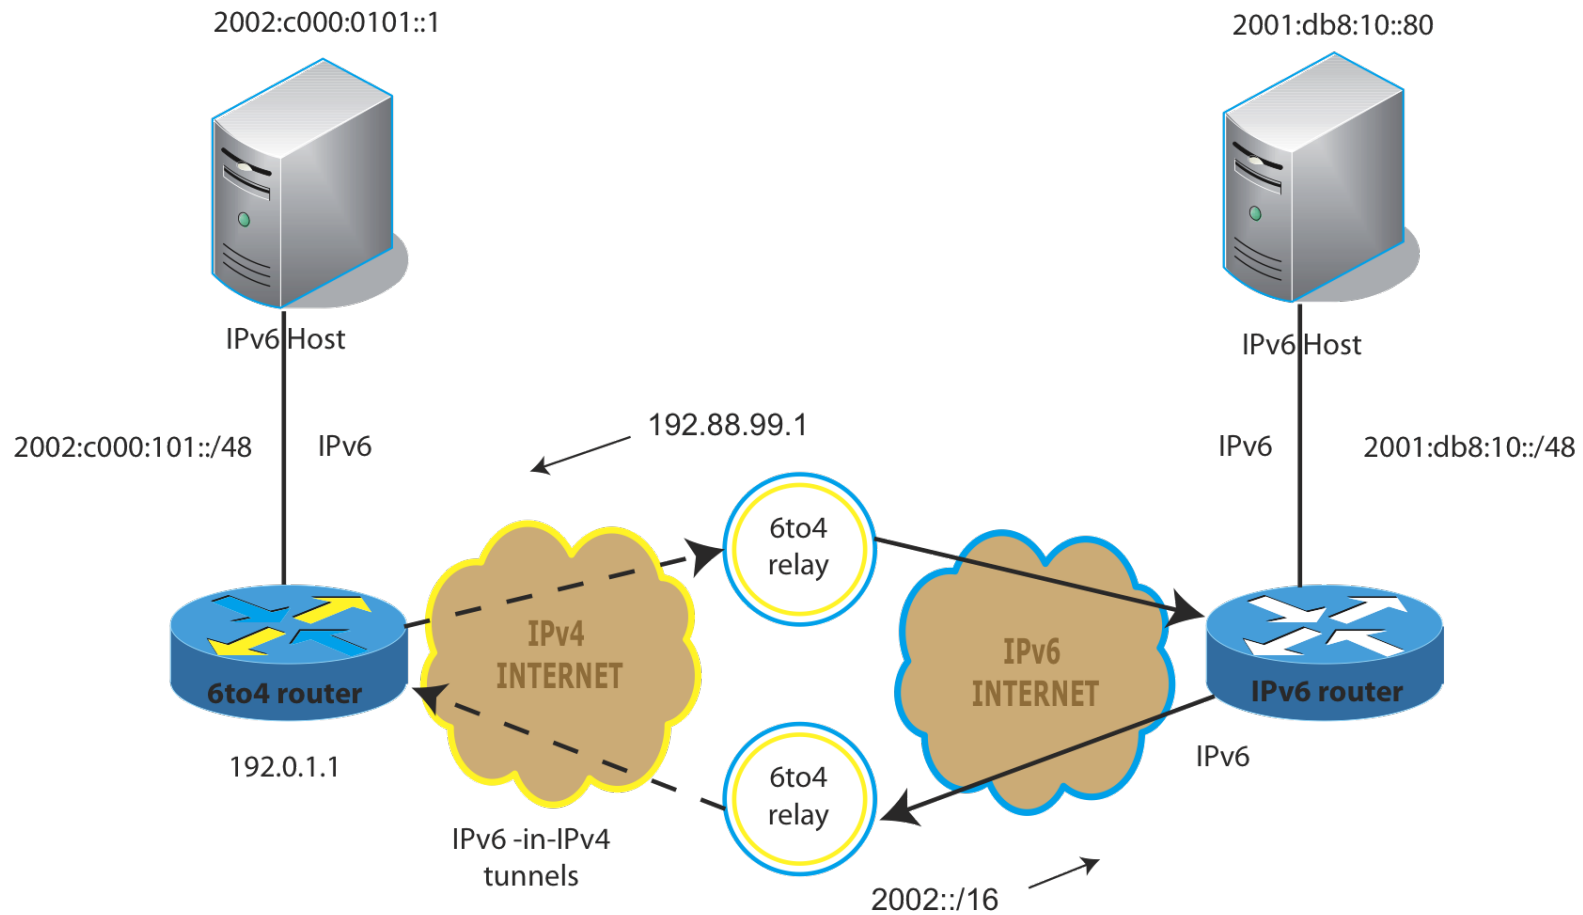 - Shows really bad connection failure rates when IPv6 literals tested (forcing IPv6 choice)
  - http://www.ietf.org/proceedings/80/slides/v6ops-1.pdf
  - It's not good…

6to4/Teredo Performance

Avg V6 Load Time Differential (secs) by V6 Type

IPv6 Connection Failure using V6 Literal

V6 Failed Connections (*)

Legend:
- 6to4 IPv6 Failure Rate (red)
- Teredo IPv6 (SYN) Failure Rate (green)
- Teredo IPv6 (ICMP) Failure Rate (blue)
- Unicast IPv6 Failure Rate (purple)

X-axis: Date (05/02, 12/02, 19/02, 26/02, 05/03, 12/03, 19/03, 26/03, 02/04)
Y-axis: % Connections (0 to 40)

6to4 highly dependent on relay

6to4 at the IETF

- In the Prague IETF in March 2011 the IETF agreed by consensus to progress two drafts:
- Moving 6to4 to Historic
  - draft-ietf-v6ops-6to4-to-historic-00
  - Means 6to4 should be disabled by default
  - Should be no new implementation of 6to4
  - Potentially no longer route 2002::/16 prefix
- Advisory on use of 6to4
  - draft-ietf-v6ops-6to4-advisory-00
  - Calls for ISPs to deploy local 6to4 relays
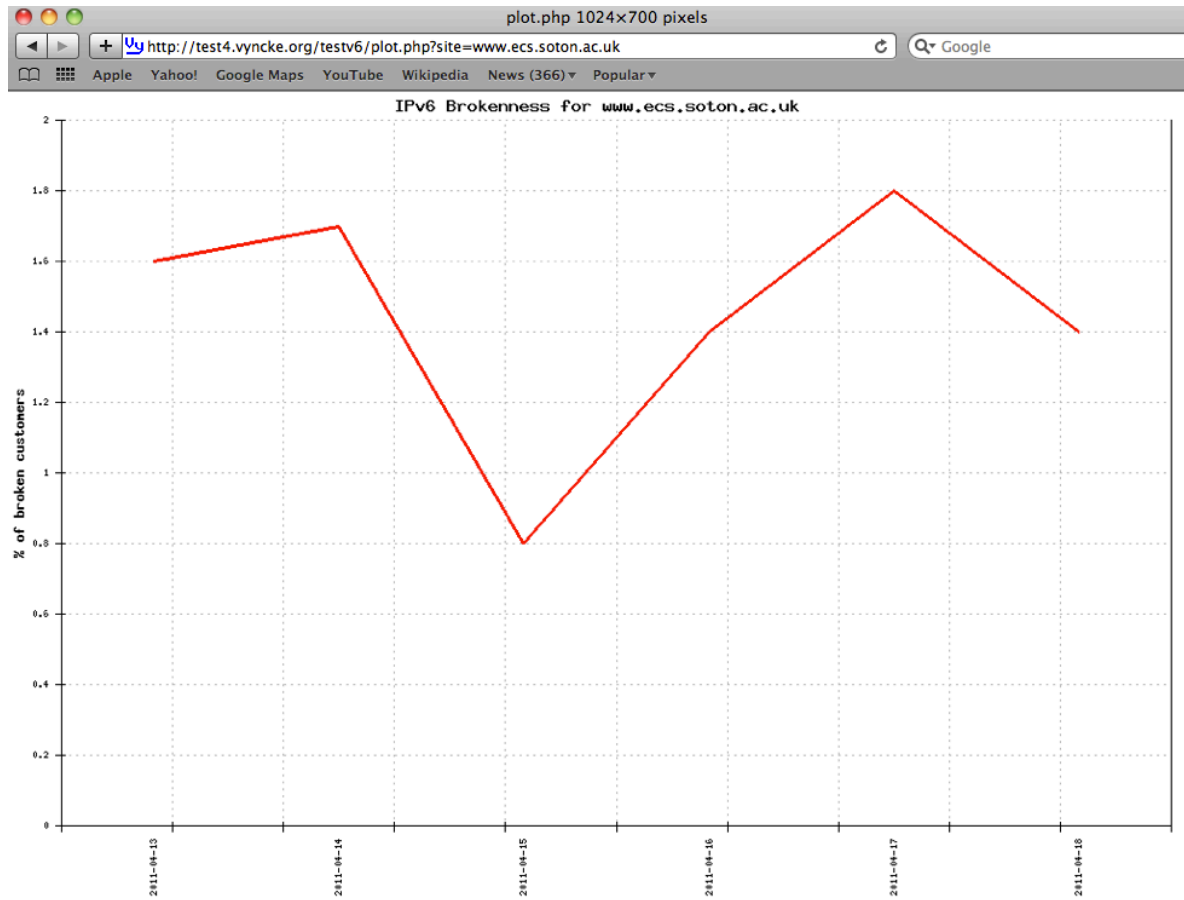 - How to make the best of 6to4 if absolutely wanted

Measurement

- Great opportunity to get good data
  - If the right tools are in place
- What could be captured on the day:
  - IPv6 traffic levels
  - Flow records (e.g. Netflow v9 to nfsen collector)
  - Application brokenness/preference
    - Reachability to IPv4-only/dual-stack/IPv6-only web server
  - PMTUD brokenness
  - IPv4/IPv6 performance comparison
  - Security monitoring (e.g. RAmond, NDPmon)

Brokenness measurements

- See next talk about tools at http://labs.apnic.net
  - We plan to use those shortly (need Google Analytics)
- We're currently testing a tool written by Eric Vynke
  - http://test4.vyncke.org/testv6
  - Just add an IFRAME to your site
  - Can view brokenness stat over time
    - Indication of potential failure rate if you dual-stack your site
    - http://test4.vyncke.org/testv6/index.php?site=www.ecs.soton.ac.uk
  - Reports User Agent strings for failed tests
- Similar to approach by Tore Anderson
  - http://ripe61.ripe.net/presentations/162-ripe61.pdf

vyncke.org brokenness

Conclusions

- We believe World IPv6 Day is a Good Thing
  - Already running IPv6 dual-stack in production
  - A glimpse into a future with Facebook, Google and others IPv6-enabled is very interesting
- Do all we can to mitigate connectivity issues
  - No problem 'affecting the experiment'
  - Detect and understand any remaining failures
- Aim to get good measurements
  - Get the tools in place now, and leave them in place
- Plan to finish IETF I-D soon as Informational
  - Would welcome any comments on it