Measuring IPv6 Deployment

Geoff Huston

APNIC R&D

August 2010

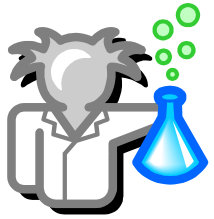

In collaboration with:

- George Michaelson, APNIC 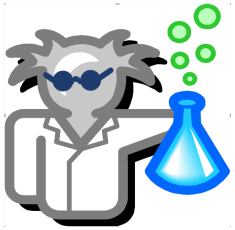
- Byron Ellacot, APNIC 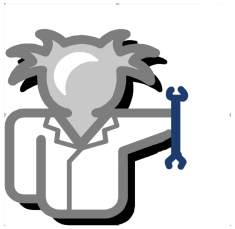
- Emile Aben, RIPE NCC 

Modelling IPv4 Address Exhaustion

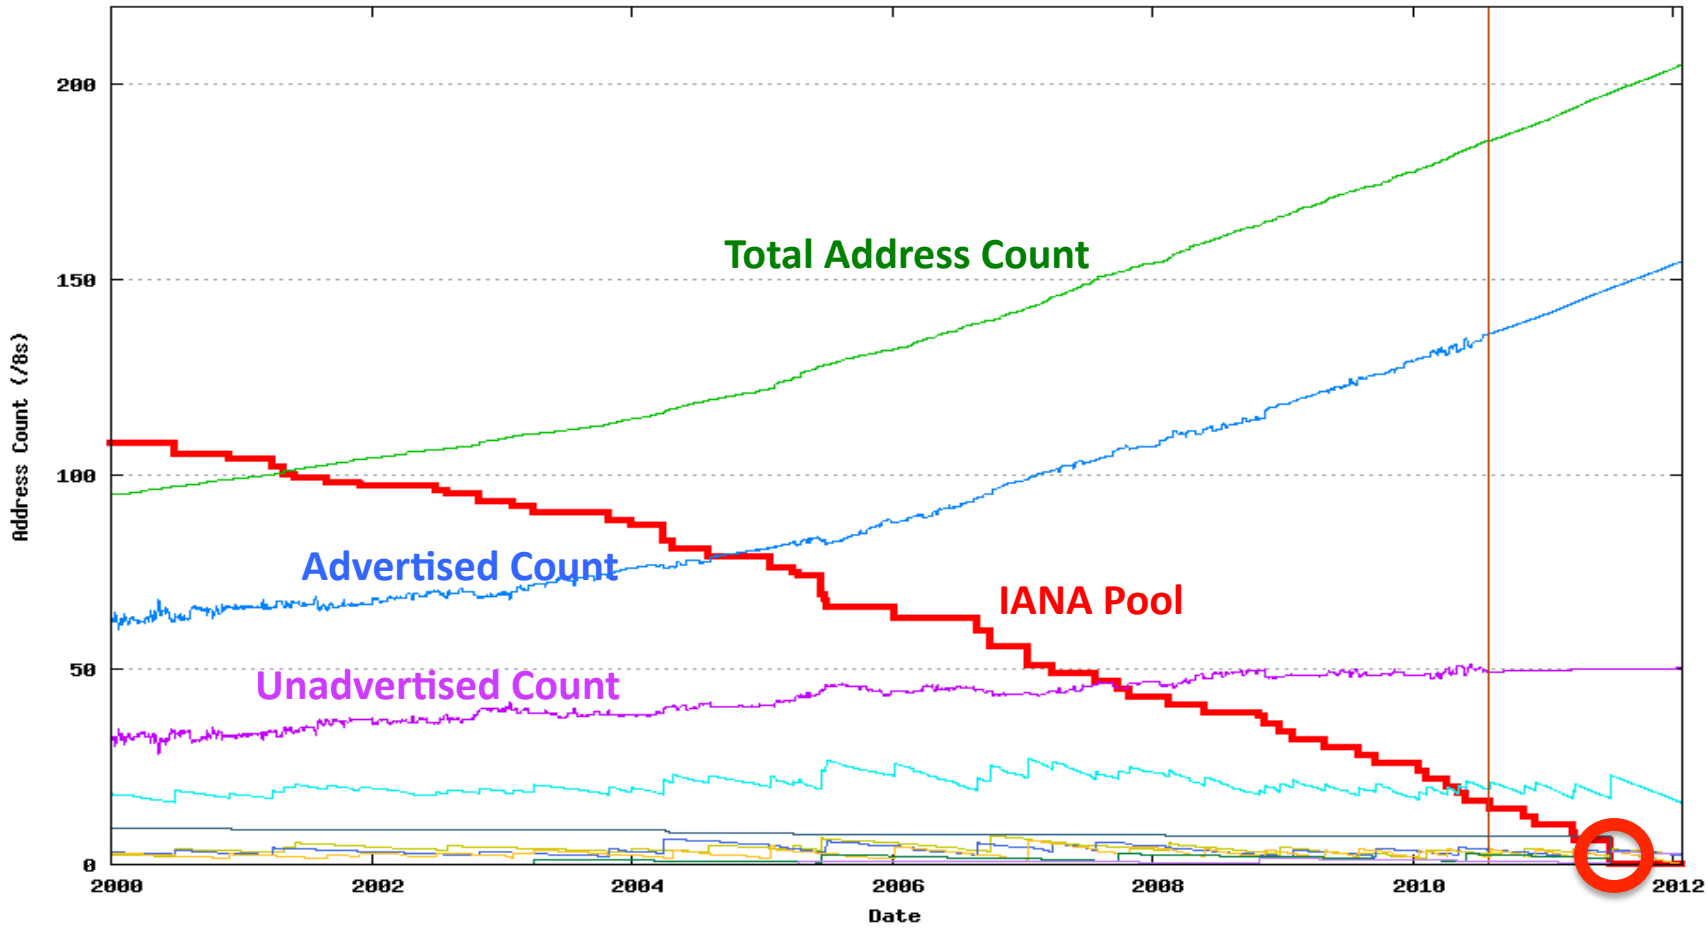

Modelling IPv4 Address Exhaustion

The model of address consumption predicts that IANA will allocate its last 'general use /8 address block on **20 June 2011**

The first RIR to exhaust its "general use" address pool will be APNIC, on **23 February 2012**

http://ipv4.potaroo.net

How certain is that date?

- As behaviours change and as policies change, the model of exhaustion changes

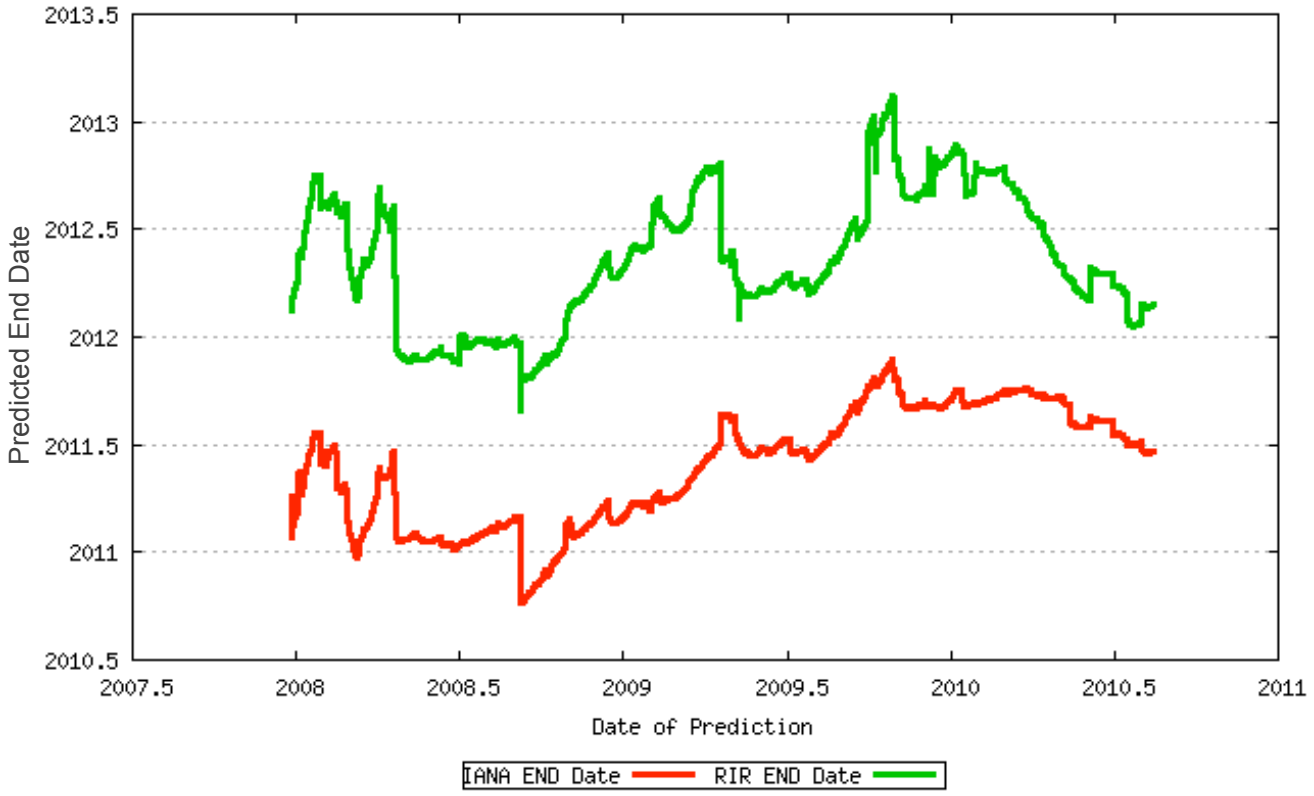

But - IPv4's Time is Running Out

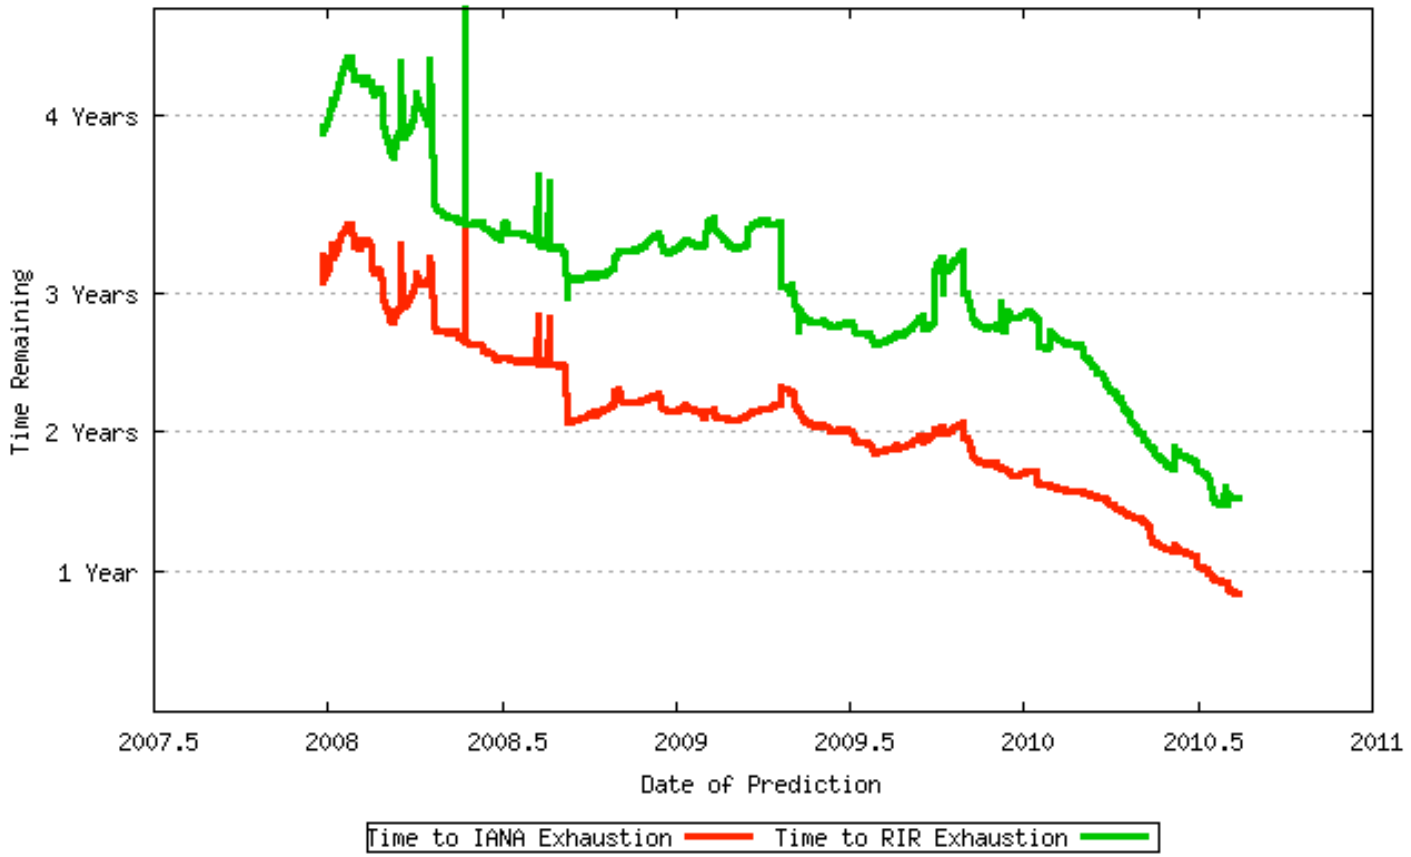

Twenty years ago…

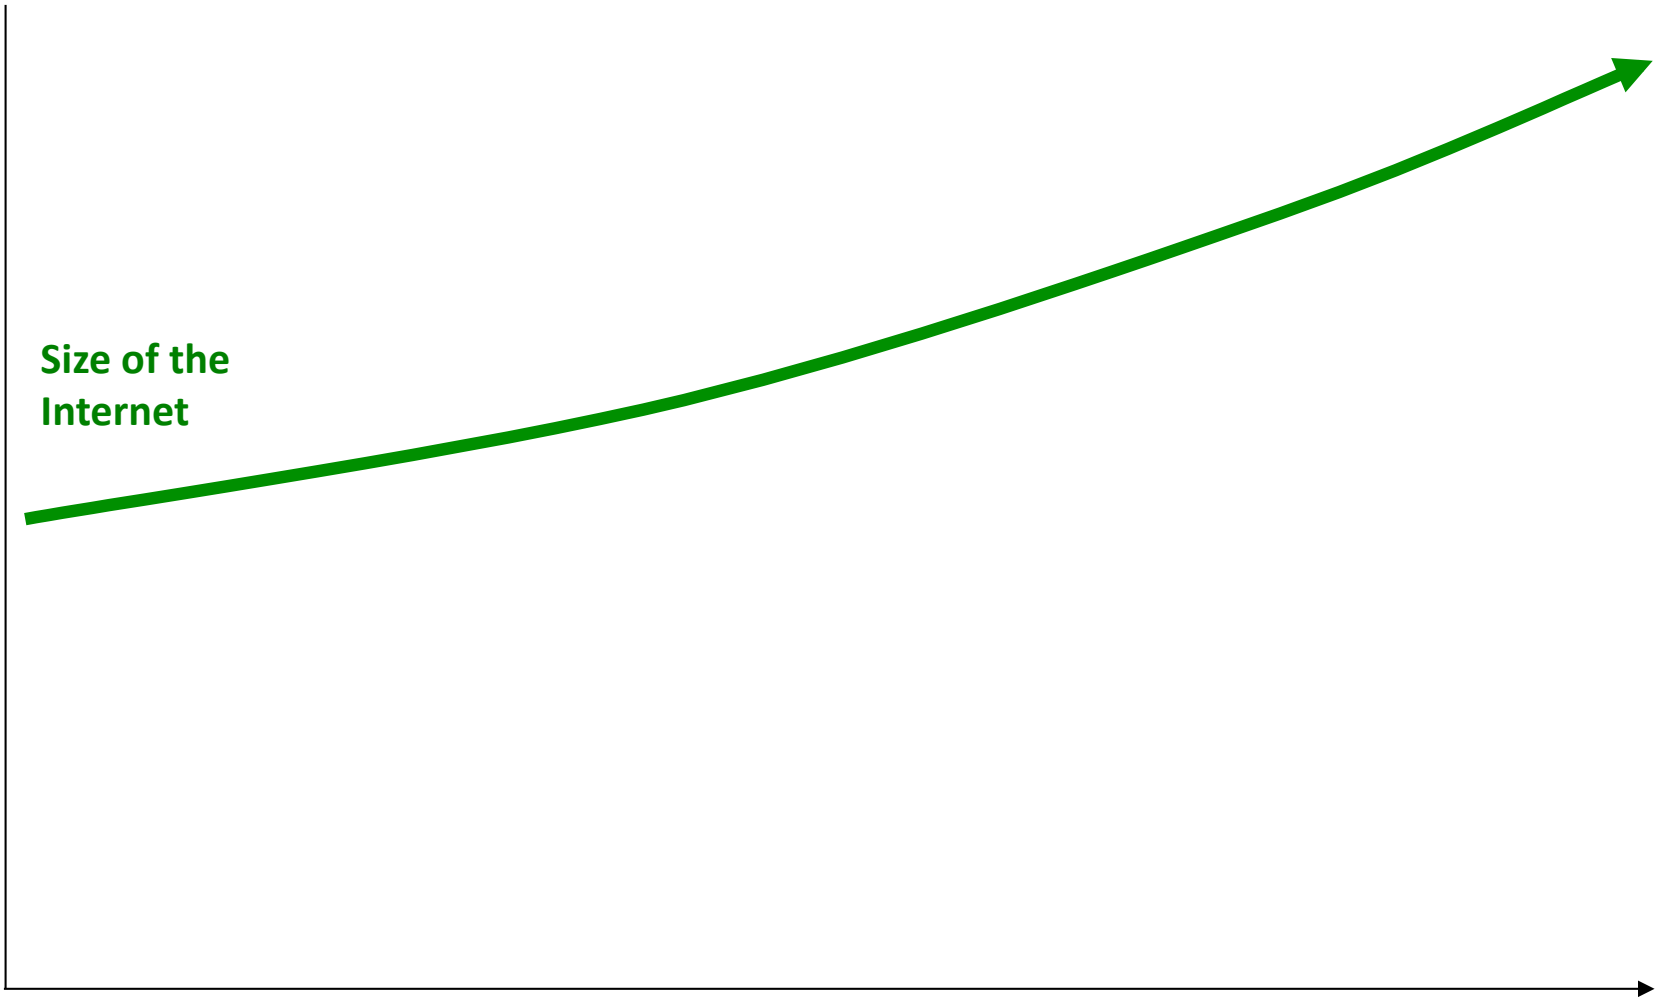

we predicted IPv4 address pool exhaustion…

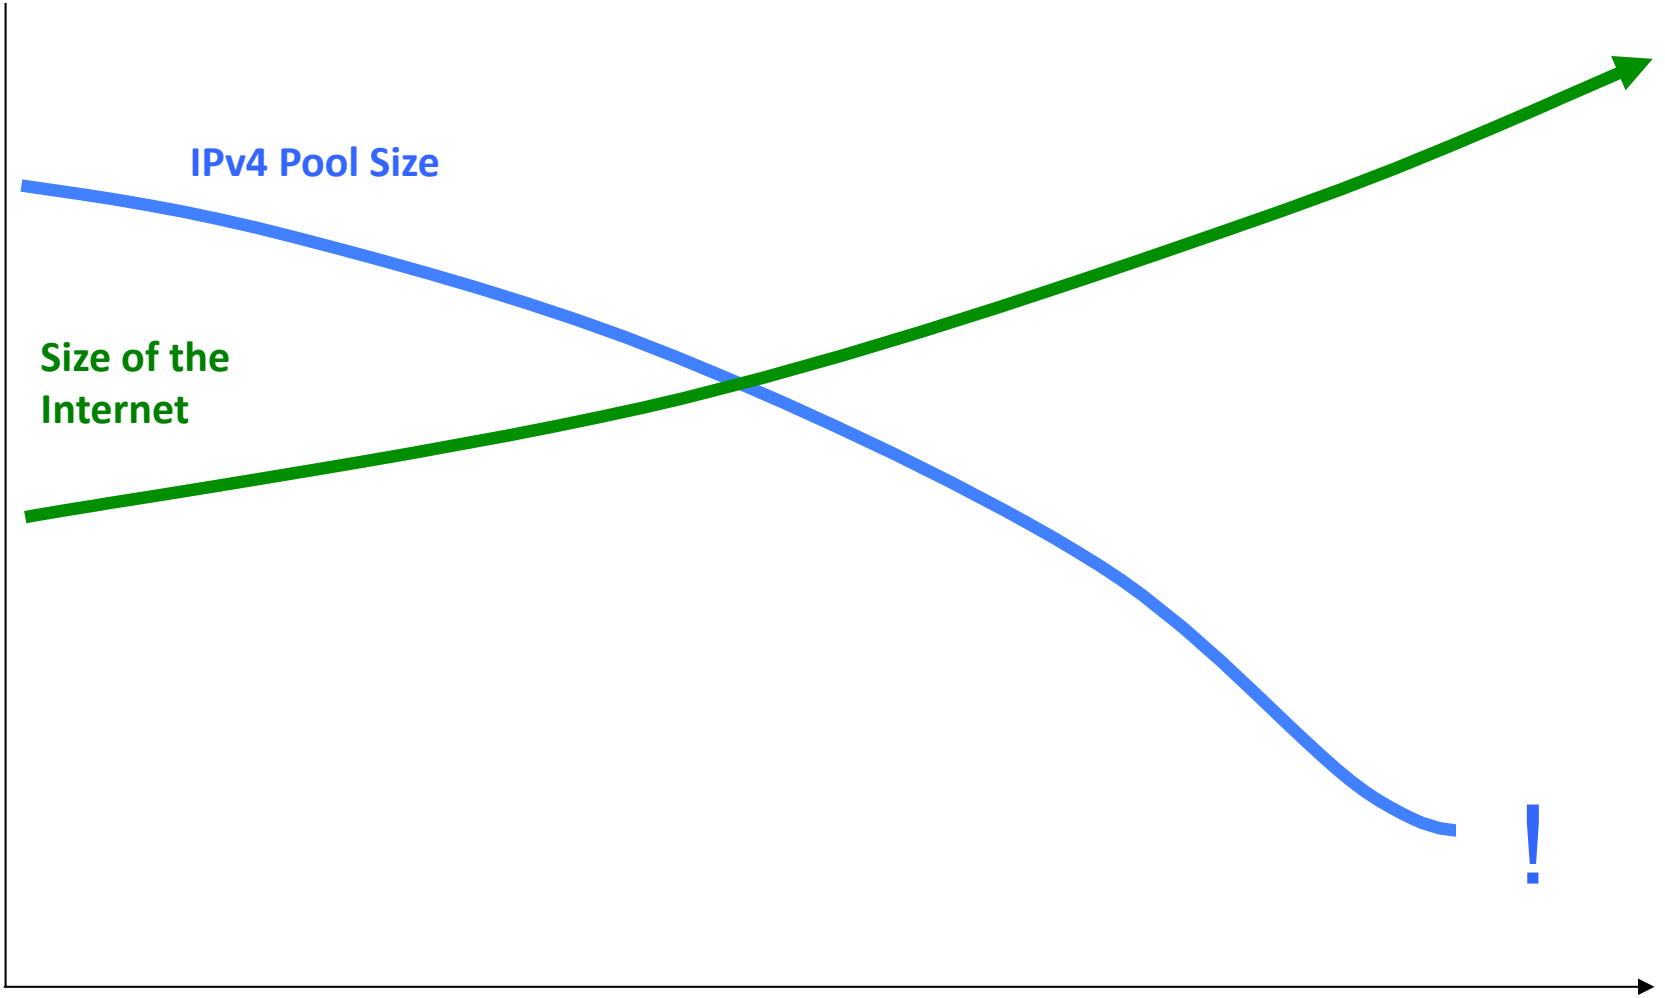

so we formed a plan…

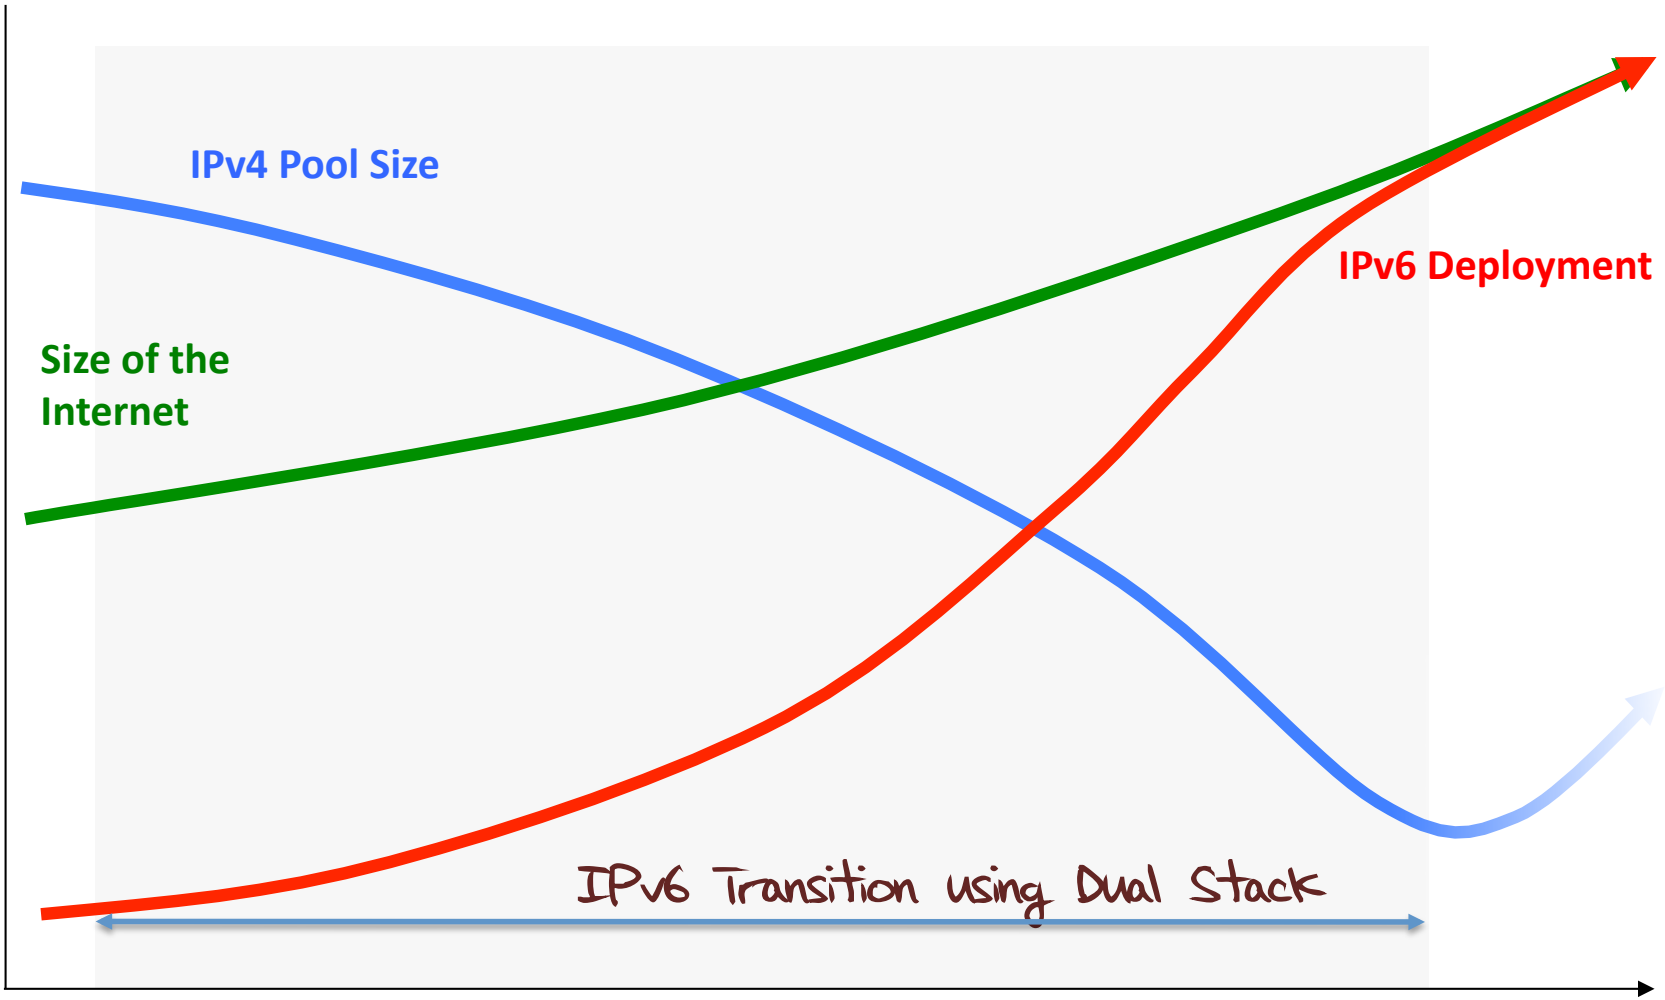

How are we going with this plan?

How are we going with this plan?

What network metrics can indicate the pace of relative deployment of IPv6?

What's the question?

Candidate questions:

- How much of the public Internet supports v6?
- How much of the public Internet runs v6?
- How quickly is the Internet becoming end-to-end v6 capable?

How long will this dual stack transition take?

Measurement Approaches

- whole of network metrics vs sample measurements
- system metrics vs component metrics
- snapshot metrics vs time series measurements

Measurement Issues

- Understanding of the nature of the environment
- The nature of the data set vs the nature of the measurement
  - Is the data set distribution heavy-tail or Gaussian?
  - Is the time series fractal or converging?
  - Is the measurement reflective of a high volume population or a small population of high volume actors?

End-to-end Service Data

- Examine IPv6 / IPv4 use from the perspective of a service delivery platform (web server)
- IPv6 is used by clients only when all the various IPv6 infrastructure components support IPv6, otherwise the client will fall back to IPv4 use
- Service metrics for IPv6 are reflective of end-to-end IPv6 capability
- Simple web log analysis approach that any dual stack web server can use

Web server stats

- Take a dual-homed web server:

  http://www.apnic.net

- Count the number of distinct IPv4 and IPv6 addresses per day
  - Not the number of web 'hits', just the ratio of the populations of distinct source addresses that access these sites, to reduce the relative impact of robots and crawlers on the data and normalize the data against different profiles of use
- Look at the v6 / v4 access ratio

Web Server V6:V4 Ratios

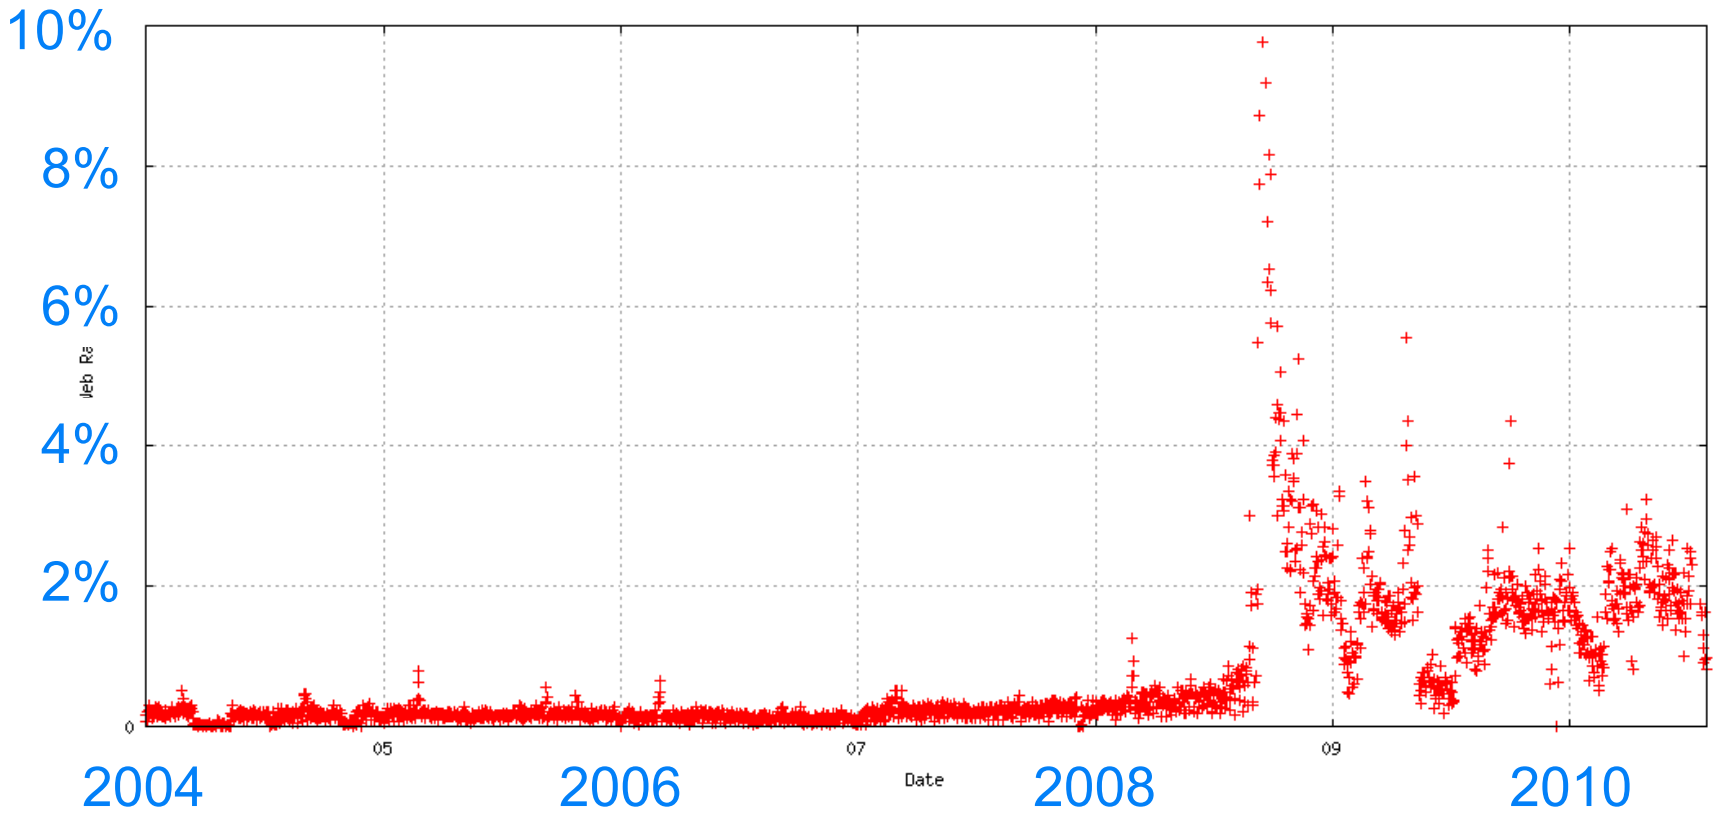

Refining the V6 Test

Instead of simply looking through the web logs for V6 / V4 ratios for web access, we can probe the client capabilities using a small script with fine-grained timers in the log

User Tests

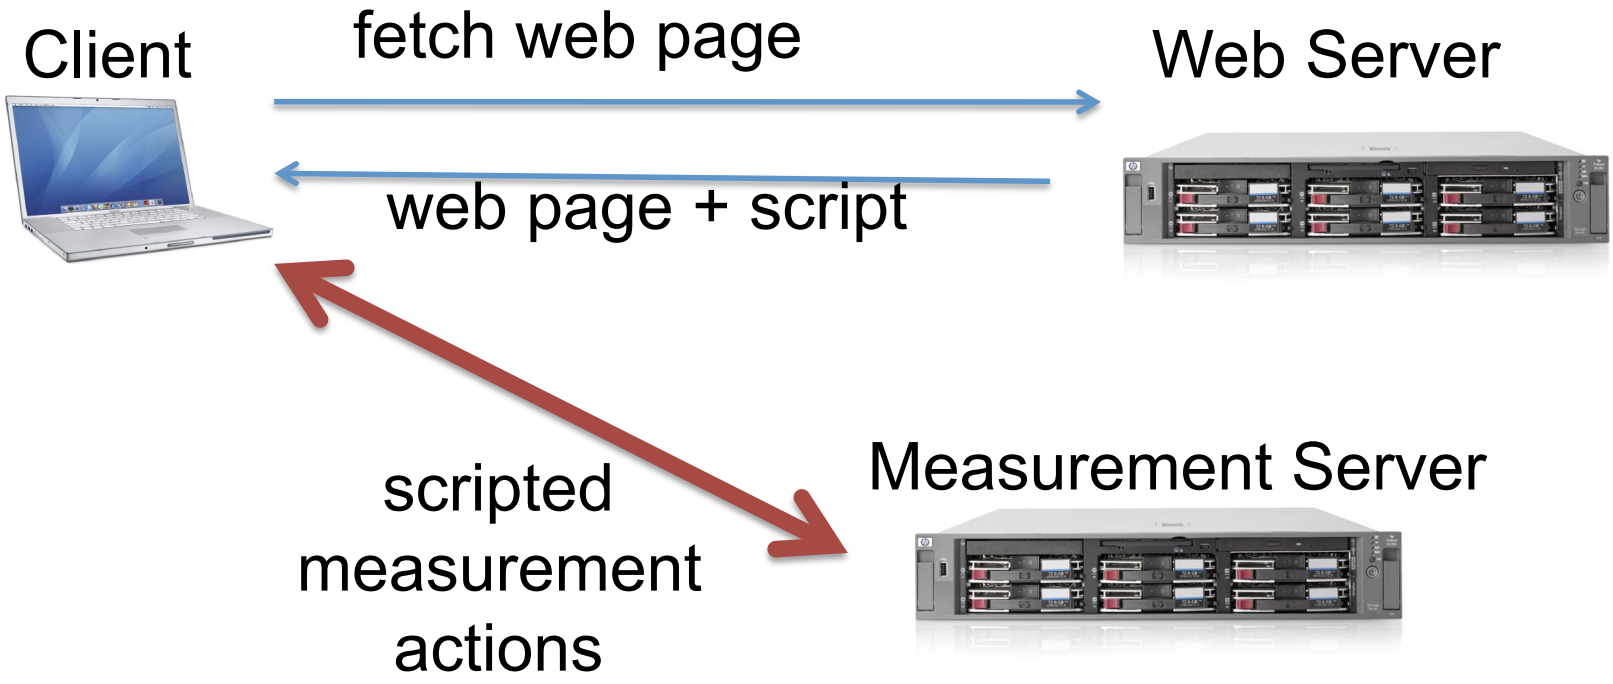
Client
fetch web page
Web Server
web page + script
scripted
measurement
actions
Measurement Server

Refining the V6 Test

Instead of simply looking through the web logs for V6 / V4 ratios for web access, we can probe the client capabilities using a small script with fine-grained timers in the log

- Test the client with 5 different retrieval tasks of a 1x1 pixel image:
  - V6 only, Dual-Stack DNS resolution
  - Dual-Stack, Dual-Stack DNS resolution
  - V4 Only, Dual-Stack DNS resolution
  - V6 Only, V6 DNS resolution
  - V4 only, V6 DNS resolution
- Use a DNS name that includes a common random number across all tests to allow server log reconstruction of the completed test set
- Take just one test result for each unique source address

Access Combinations

V4	V6	Dual	Node Type
✔	✖	V4	V4-Only
✖	✔	V6	V6-Only
✔	✔	V6	V6-Preferred
✔	✔	V4	V6-Capable (V4-Preferred)
✔	✖	✖	Dual-Stack Loss

IPv6: "could" vs "will"

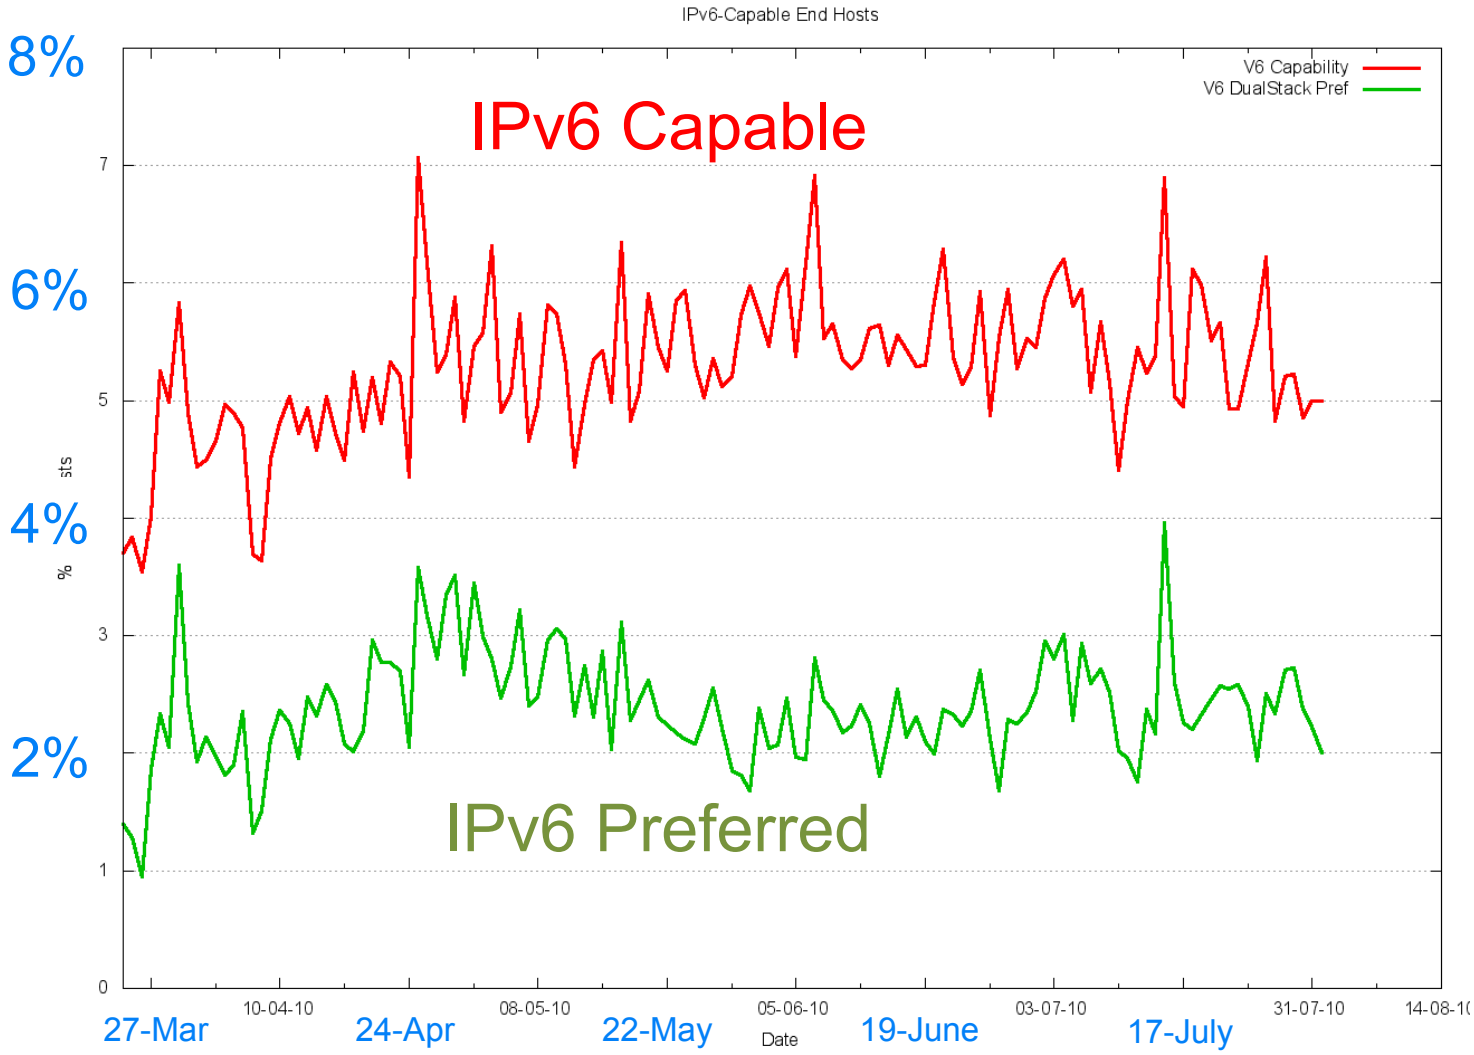

Where are we with IPv6?

The ‘size’ of the IPv6 deployment in terms of end-to-end host IPv6 preference is around 2% of the total number of Internet end hosts at present

However a further 3% of hosts can use IPv6, even though they prefer IPv4 in dual stack mode

How much V6 is "out there"?

Around 5.5% of end hosts are capable of reaching IPv6 only service points

This is more than double the number of hosts who expose their V6 capability by V6 preference in dual stack

Some V6-only hosts already?

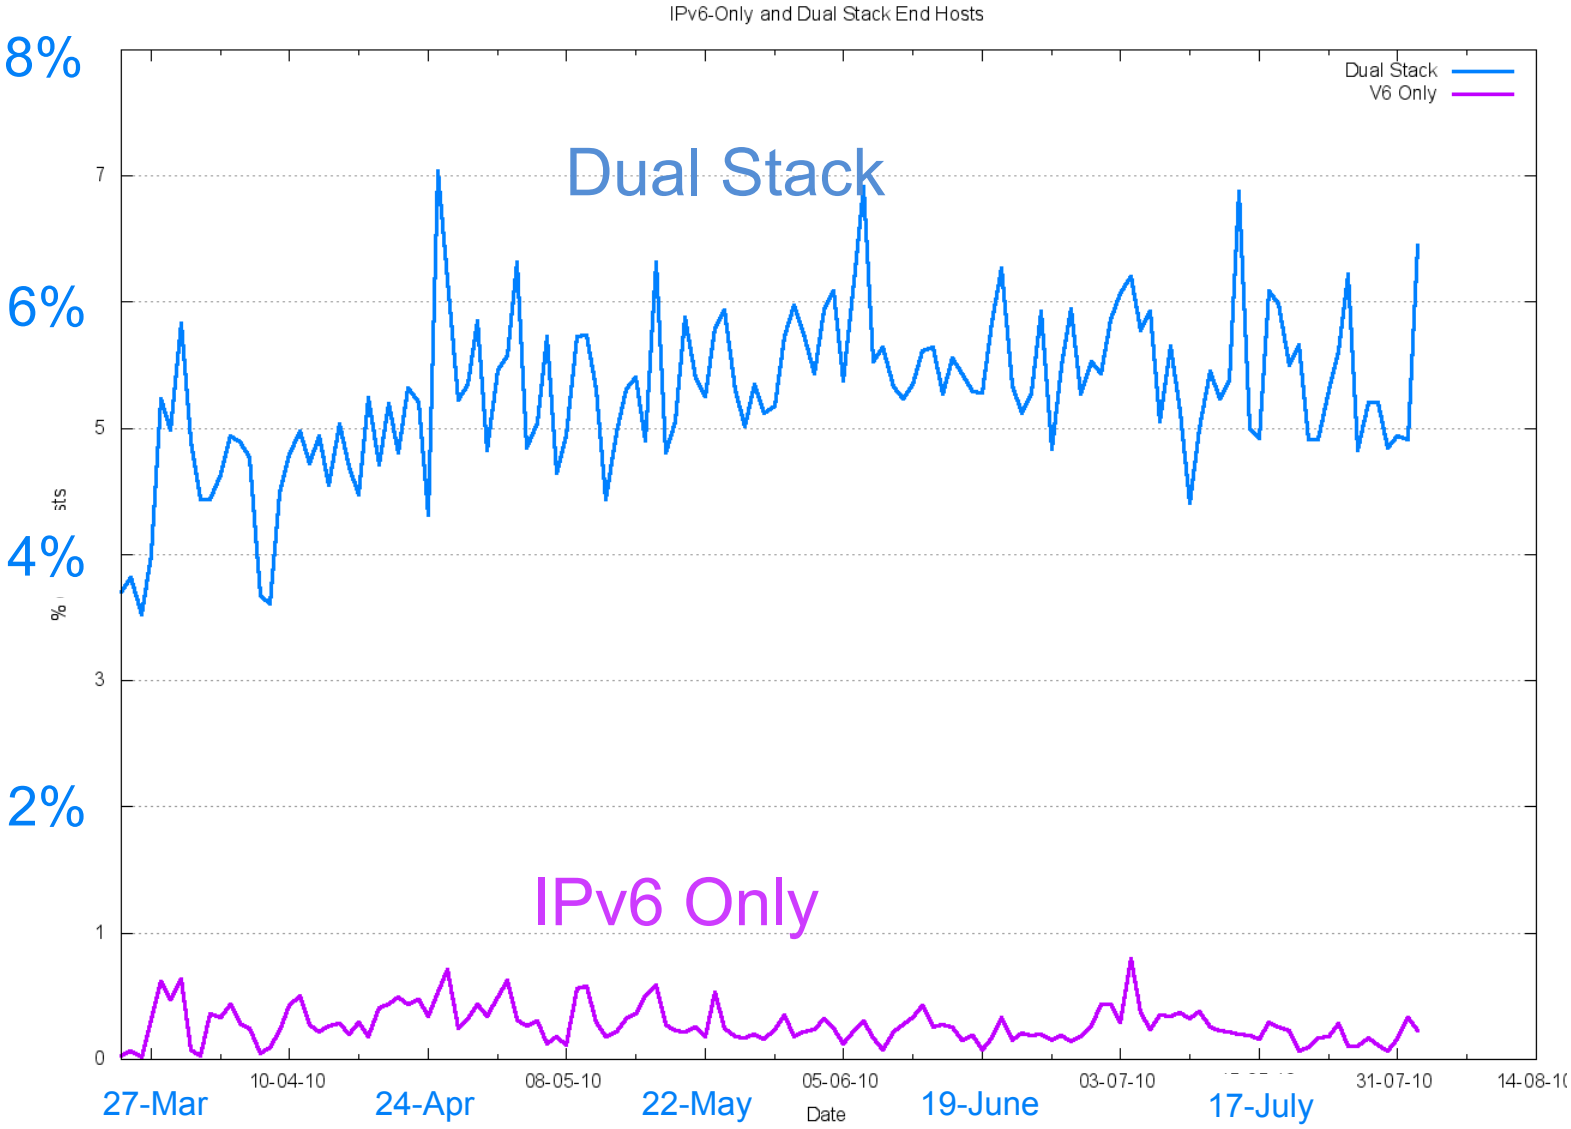

Access by V6 Address Type

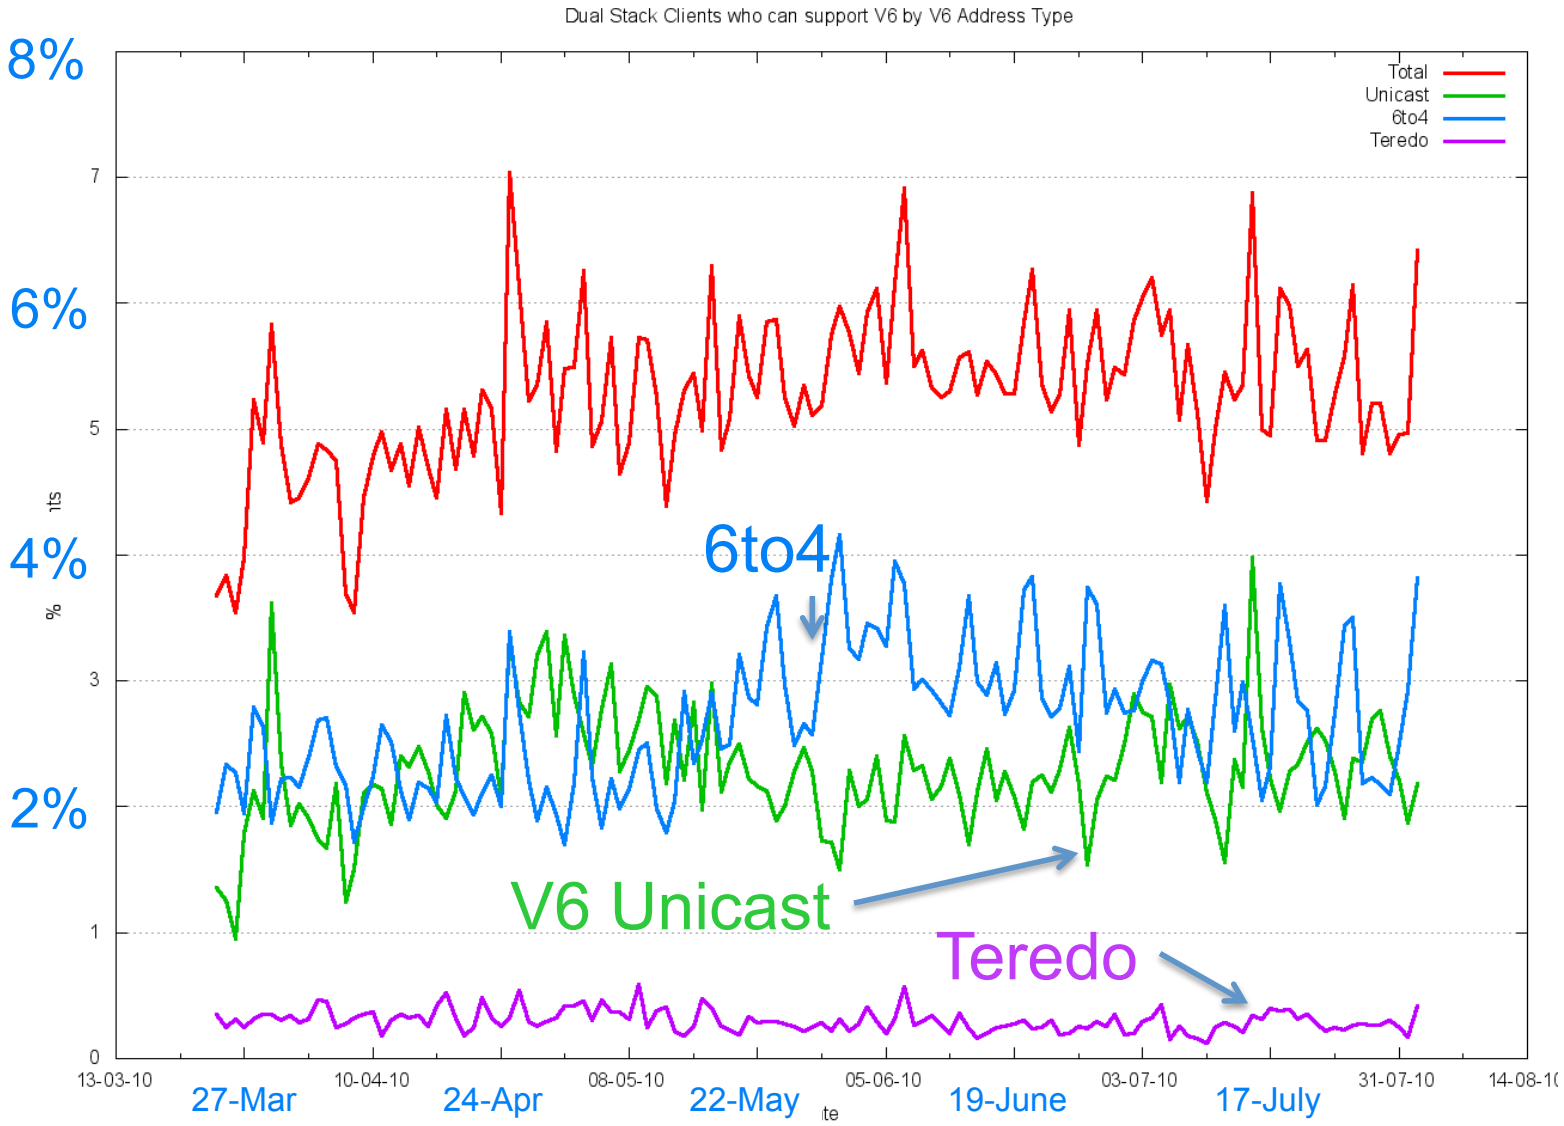

Dual-Stack, V6 Preferred by Address Type

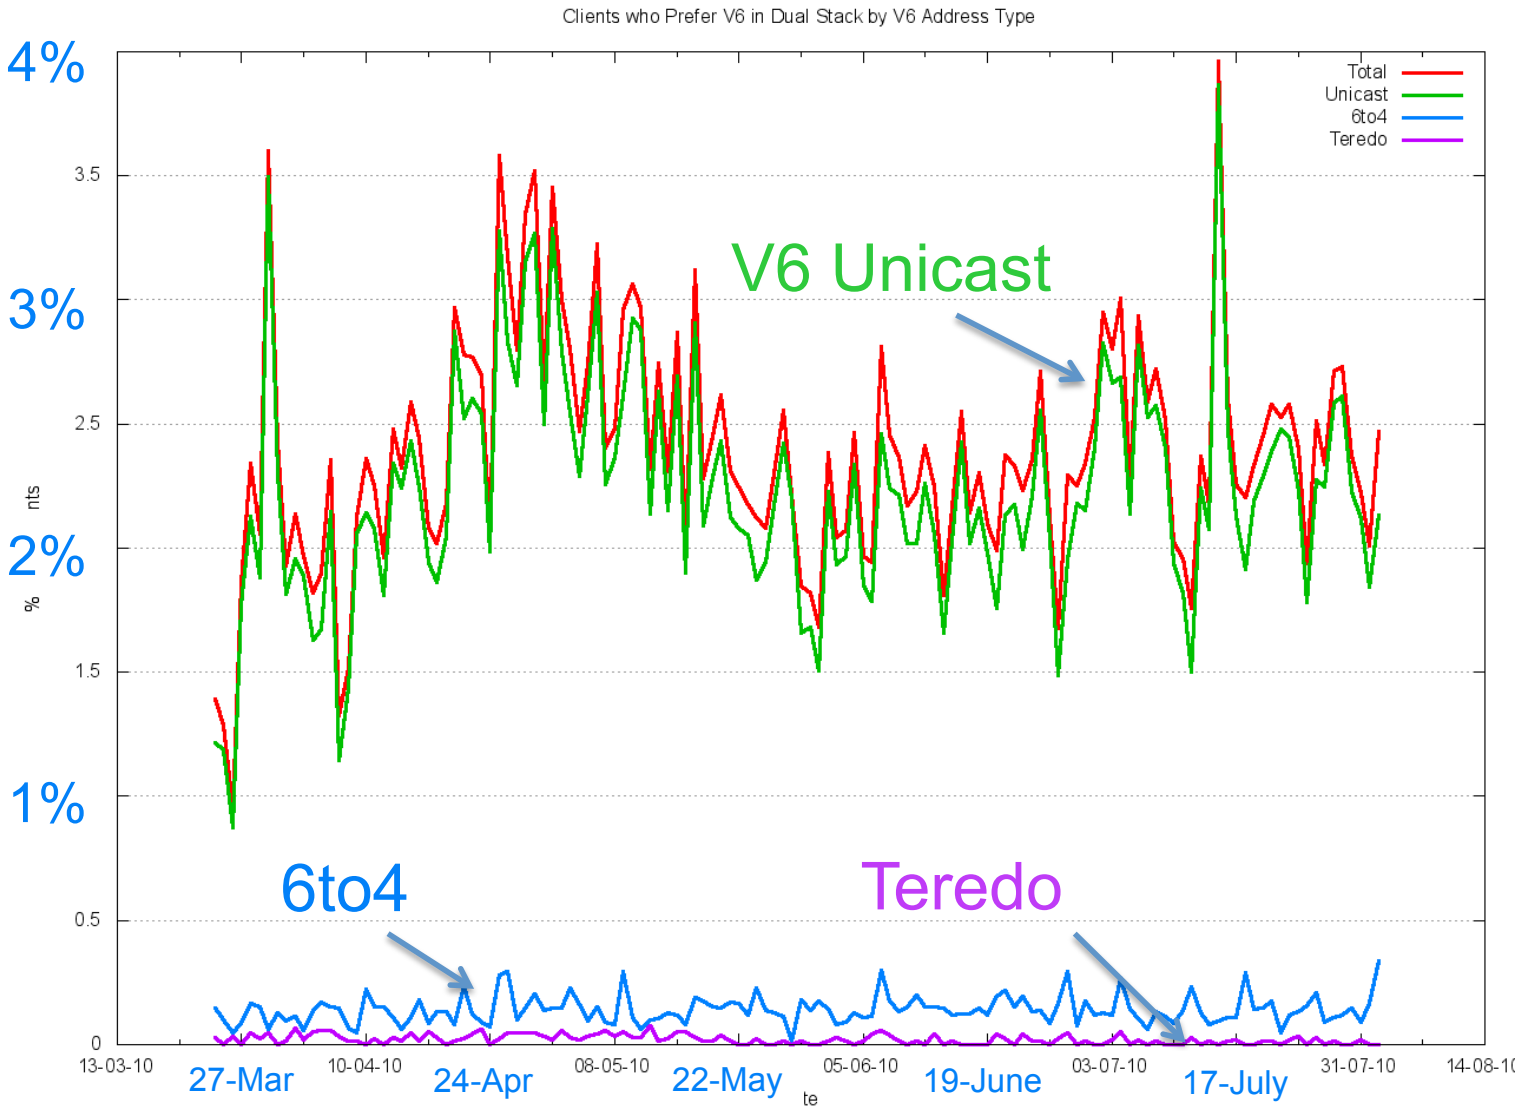

Dual-Stack, V4 Preferred by Address Type

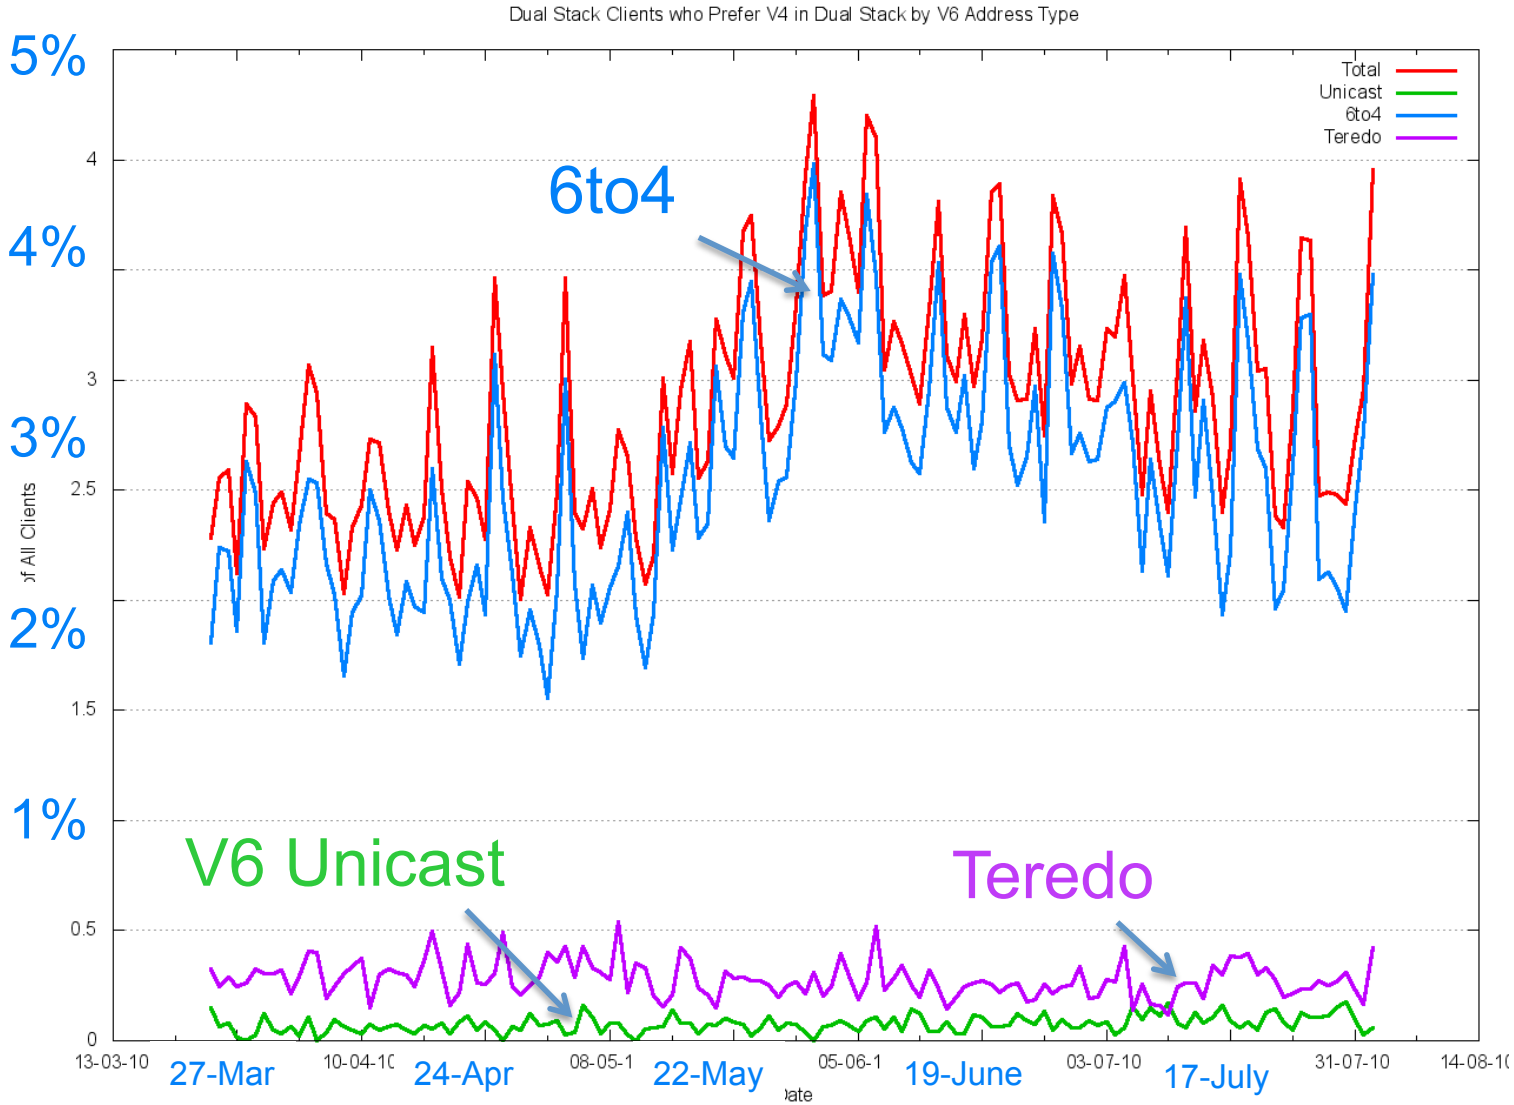

Unicast vs Tunneled

- Most hosts with unicast IPv6 generally prefer V6 in a dual stack scenario
- Hosts with auto-tunnel capability (typically Windows Vista and 7 systems) appear to generally prefer V4 in a dual stack scenario when configured to use 6to4 or Teredo auto-tunneling to access the V6 network

Teredo vs 6to4

- What we see:
  - 3% of hosts use 6to4 (native V4, auto-tunnel)
  - 0.3% of hosts use Teredo (NAT V4, auto-tunnel)
- Aren't there more hosts behind v4 NATs than hosts on native v4?
- Therefore, shouldn't we see more Teredo than 6to4 in auto-tunnel tests?
- Why is the level of Teredo usage so much smaller than 6to4?
  - could it be due to extensive use of highly restrictive filters in consumer grade IPv4 NATs?
  - or widespread disabling of Teredo in end systems?
  - or …?

Who uses 6to4?

- 6to4 auto-tunnelling access counts appear to peak on weekends
- Do corporate environments today still rely heavily on V4 NATs with V6 auto-tunnelling disabled on end systems and through the firewalls?

Performance and Tunnels

- Tunnelling can extend the packet path:
  - addition of a tunnel relay between the source and destination
  - in the case of 6to4 this is asymmetric, potentially lengthening the transit path

6to4 Packet Path

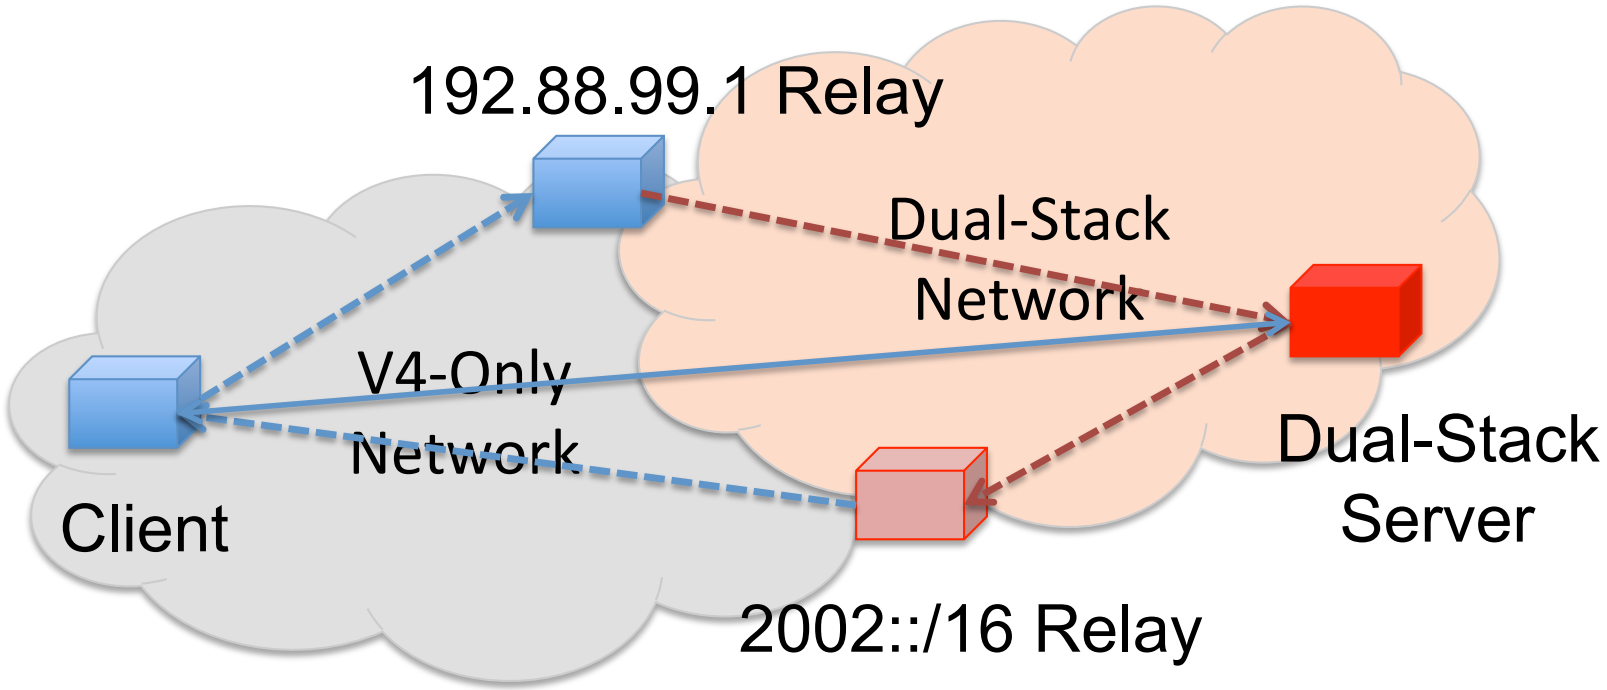

Partial Mitigation of 6to4 Packet Path

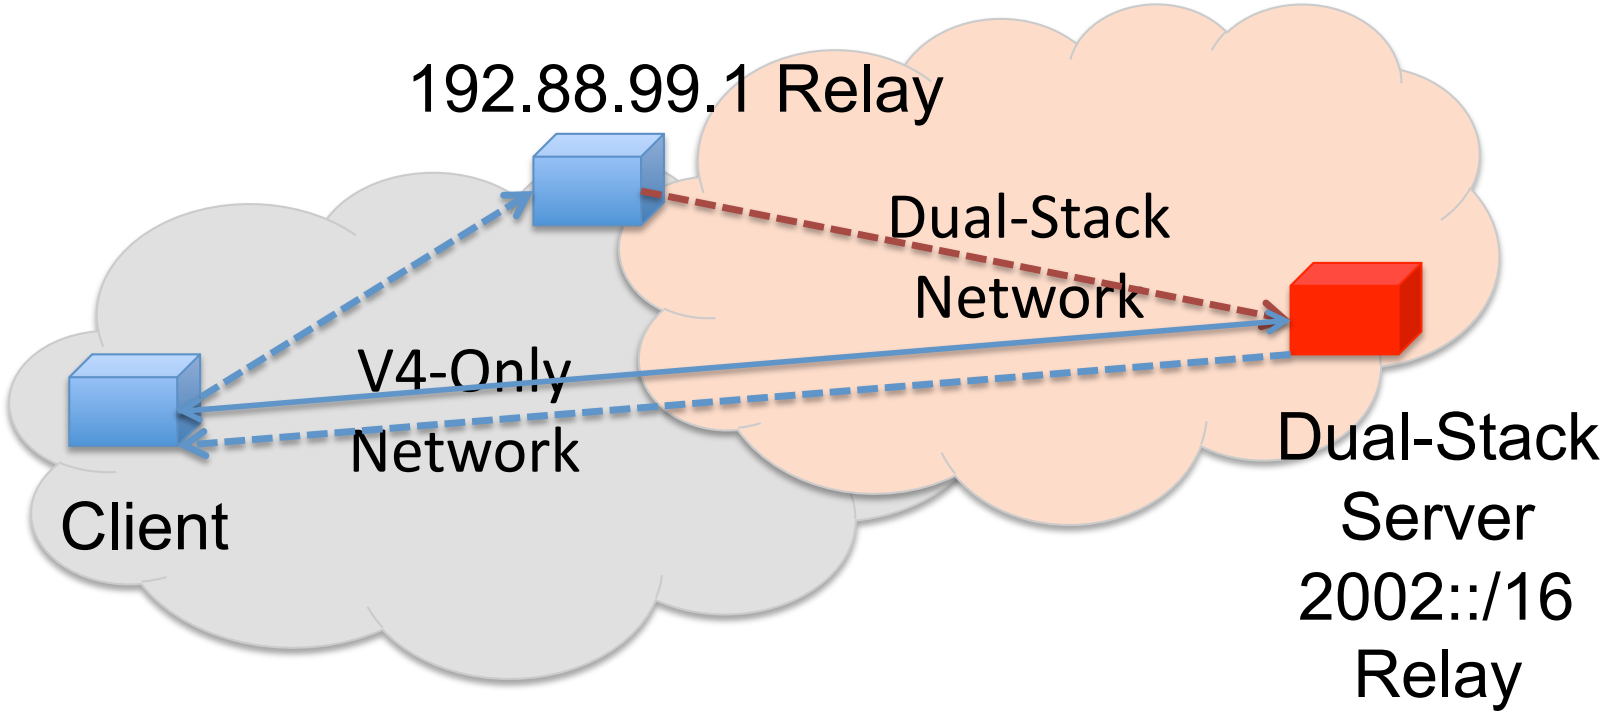

Performance and Tunnels

- Teredo can add a further performance penalty in the form of state setup between the Teredo relay and the client

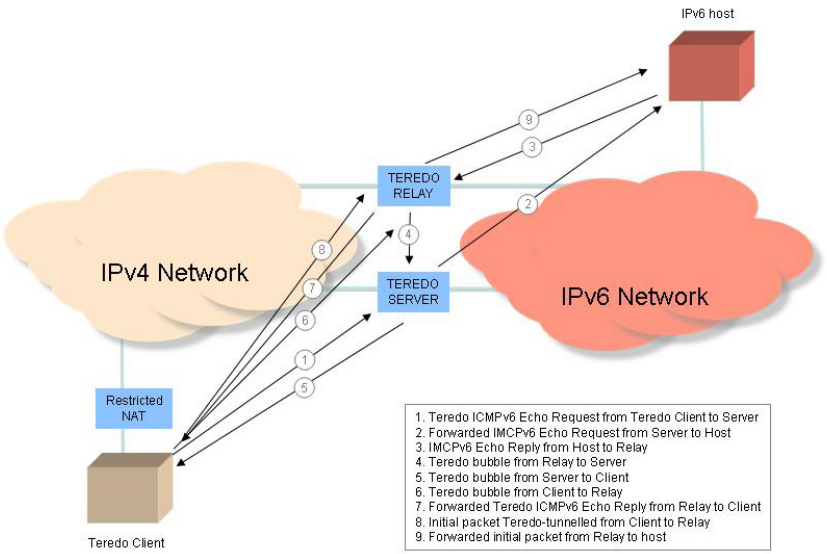

Performance and Tunnels

- How big a performance penalty do we see for V6 via tunnels?

What are we measuring?

- The script uses wildcard DNS entries and random-value DNS names to ensure that client-side web caching and DNS caching is avoided
- The server records the time of script delivery to the client and the time of delivery of each of the test objects – the difference is the retrieval time
- The retrieval time includes DNS and TCP protocol overheads
- The sequence of objects in the script is constant: the order is V6, Dual Stack, then V4
- V4 retrieval time is the benchmark against which the V6 retrieval time is measured

Performance and Tunnels

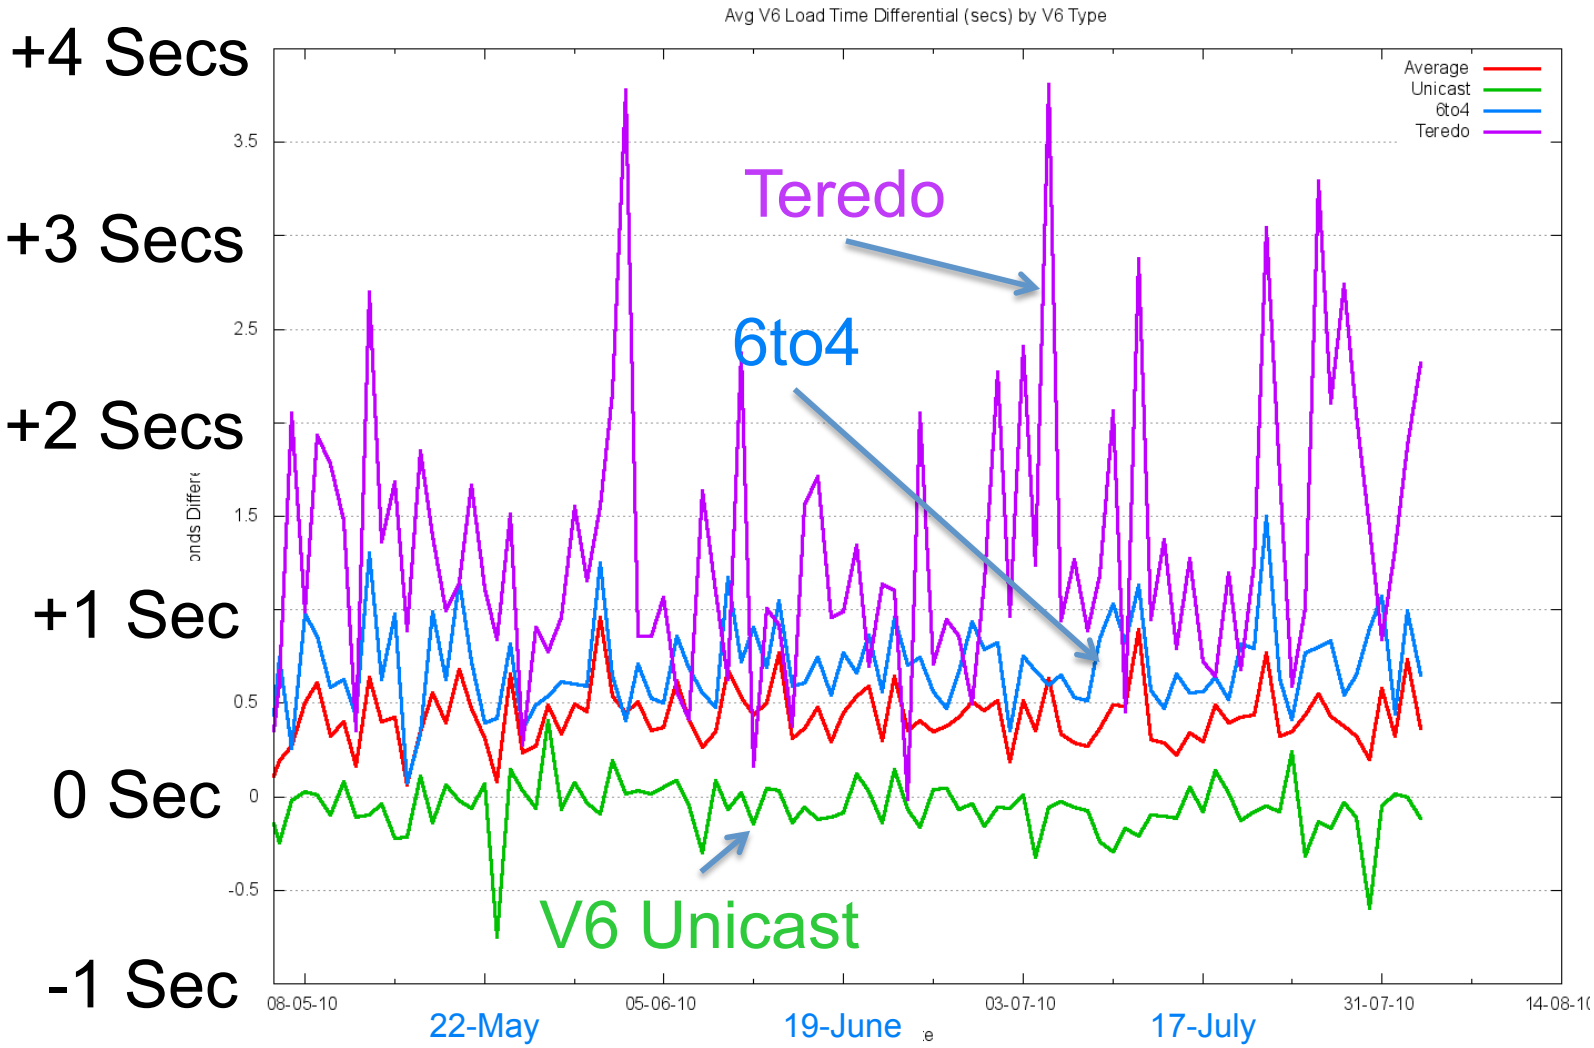

Comparison with RIPE data

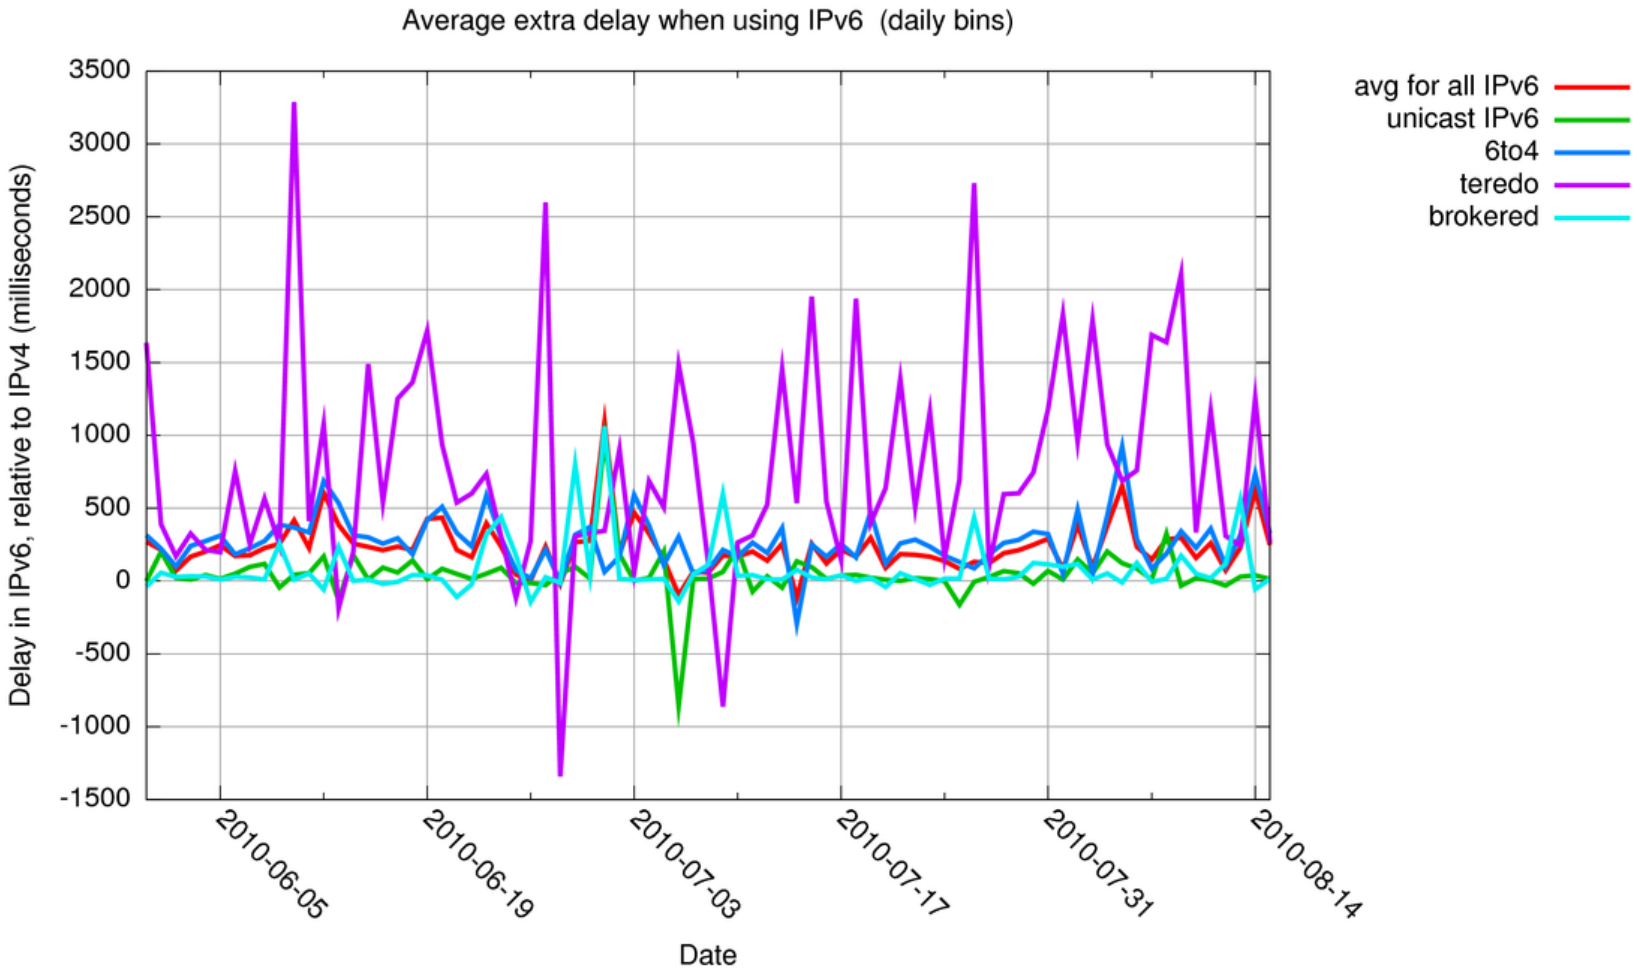

http://albatross.ripe.net/v6-clientresolver/site_ncc/
Thanks to Emile Aben, RIPE NCC

V6 Relative Performance

- Retrieval times for native V6 are the same as V4
  - (or even a little better on average, which could be due to client-side sequencing of requests provided in the script)

6to4 Relative Performance

- Auto-tunnel 6to4 adds an average of 0.6 seconds to the retrieval time
  - note this is one-way (as the server has a local 6to4 relay for the response traffic, so the 6to4 response path is the same as the V4 path)
  - that's a very long transit time if this is just added transit time
  - so there may be a congestion load delay added in here
- Are outbound 192.88.99.1 6to4 relays so sparsely deployed and so heavily overloaded that the average one way additional delay is ~600ms?
  - RIPE NCC measure a ~300ms 6to4 delay over V4
  - The higher relative delay in 6to4 could be attributed to sparse deployment of 6to4 relays in the AP region

Teredo Relative Performance

- Auto-tunnel Teredo V6 adds an average of 1 – 2 seconds to the retrieval time
  - that's a really long additional delay!
  - what is causing this significantly higher performance penalty?

The additional delay here could be caused by the combination of the Teredo setup phase, and the server's use of a remote 2001::/32 relay server

For the next phase of this work we are looking at equipping the measurement server with a local 2001::/32 Teredo relay server to eliminate the relay server hop on the reverse path to the client

V6 Performance

- Unicast V6 appears to be as fast as V4
- Auto-tunnel V6 does attract some performance overheads
  - these are strongly context dependant
  - widespread deployment of 6to4 relays and Teredo relays and servers will help
  - Dual Stack servers may want to consider using local 6to4 and Teredo relays to improve reverse path performance for auto-tunnelling clients

Dual Stack Failure

How many clients retrieve the V4 only object but DON’T retrieve the Dual Stack objects?

i.e. how many clients exhibit “Dual Stack Failure”?

Dual Stack Failure Rate

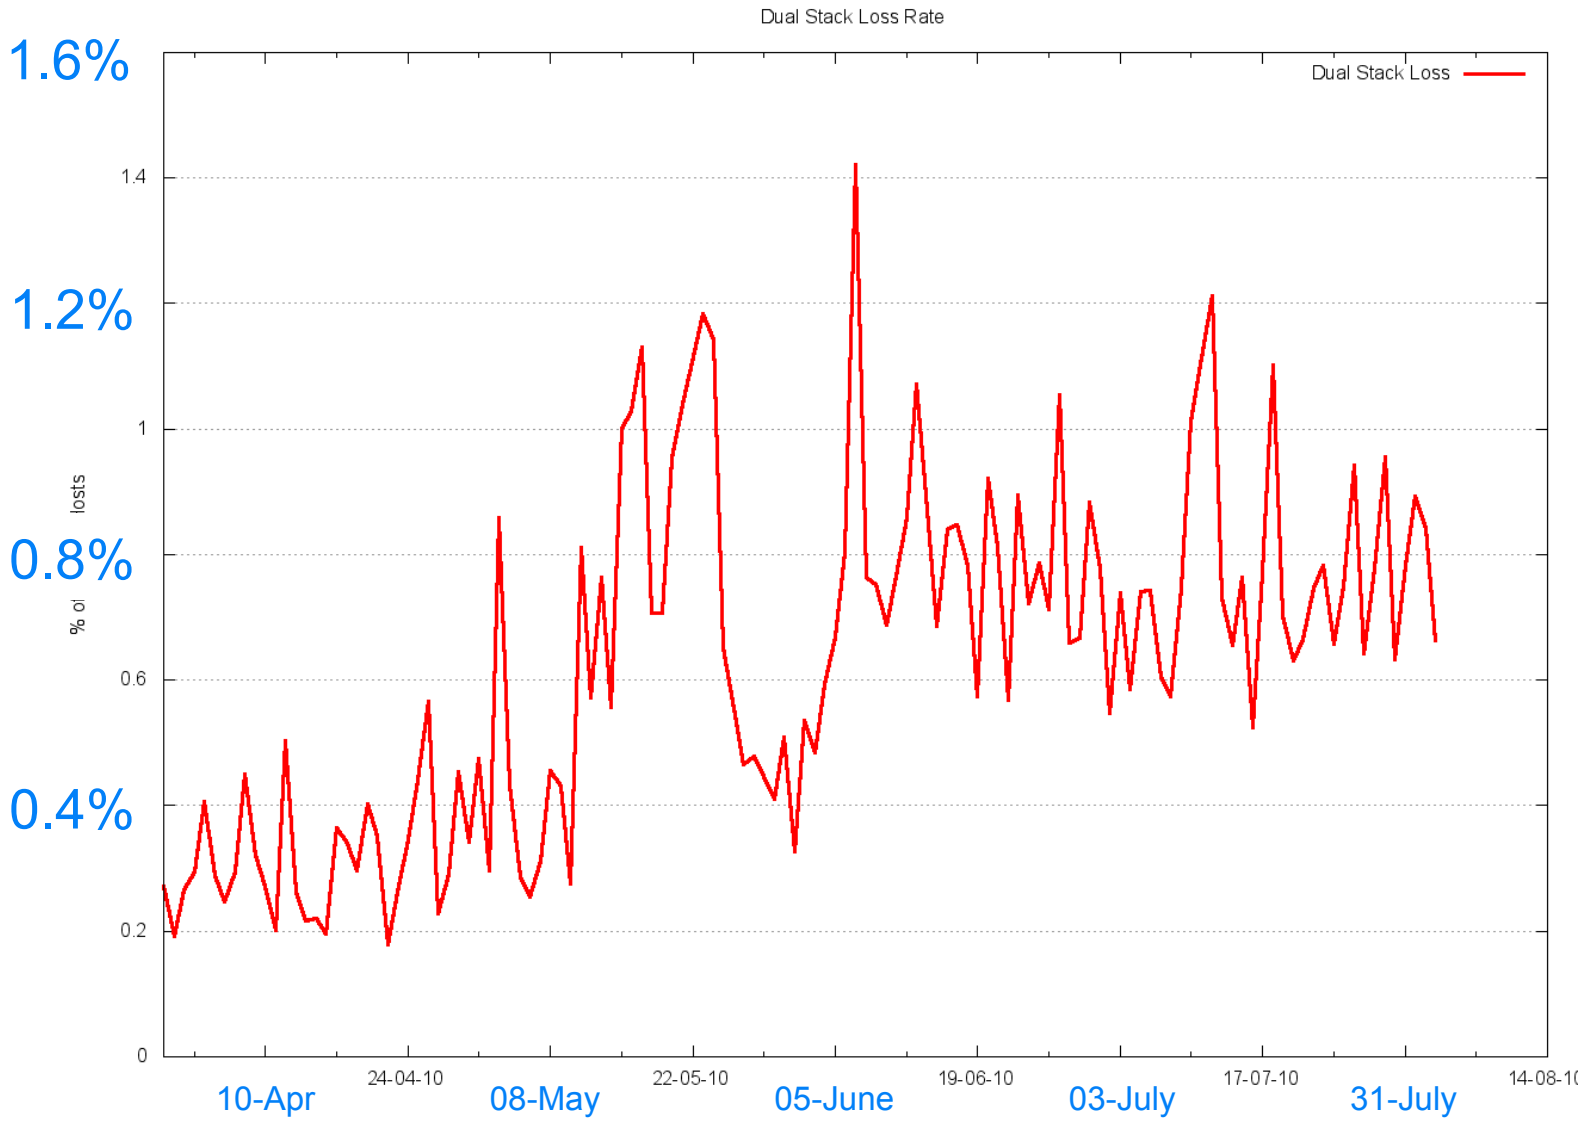

Dual Stack Failure - Windows

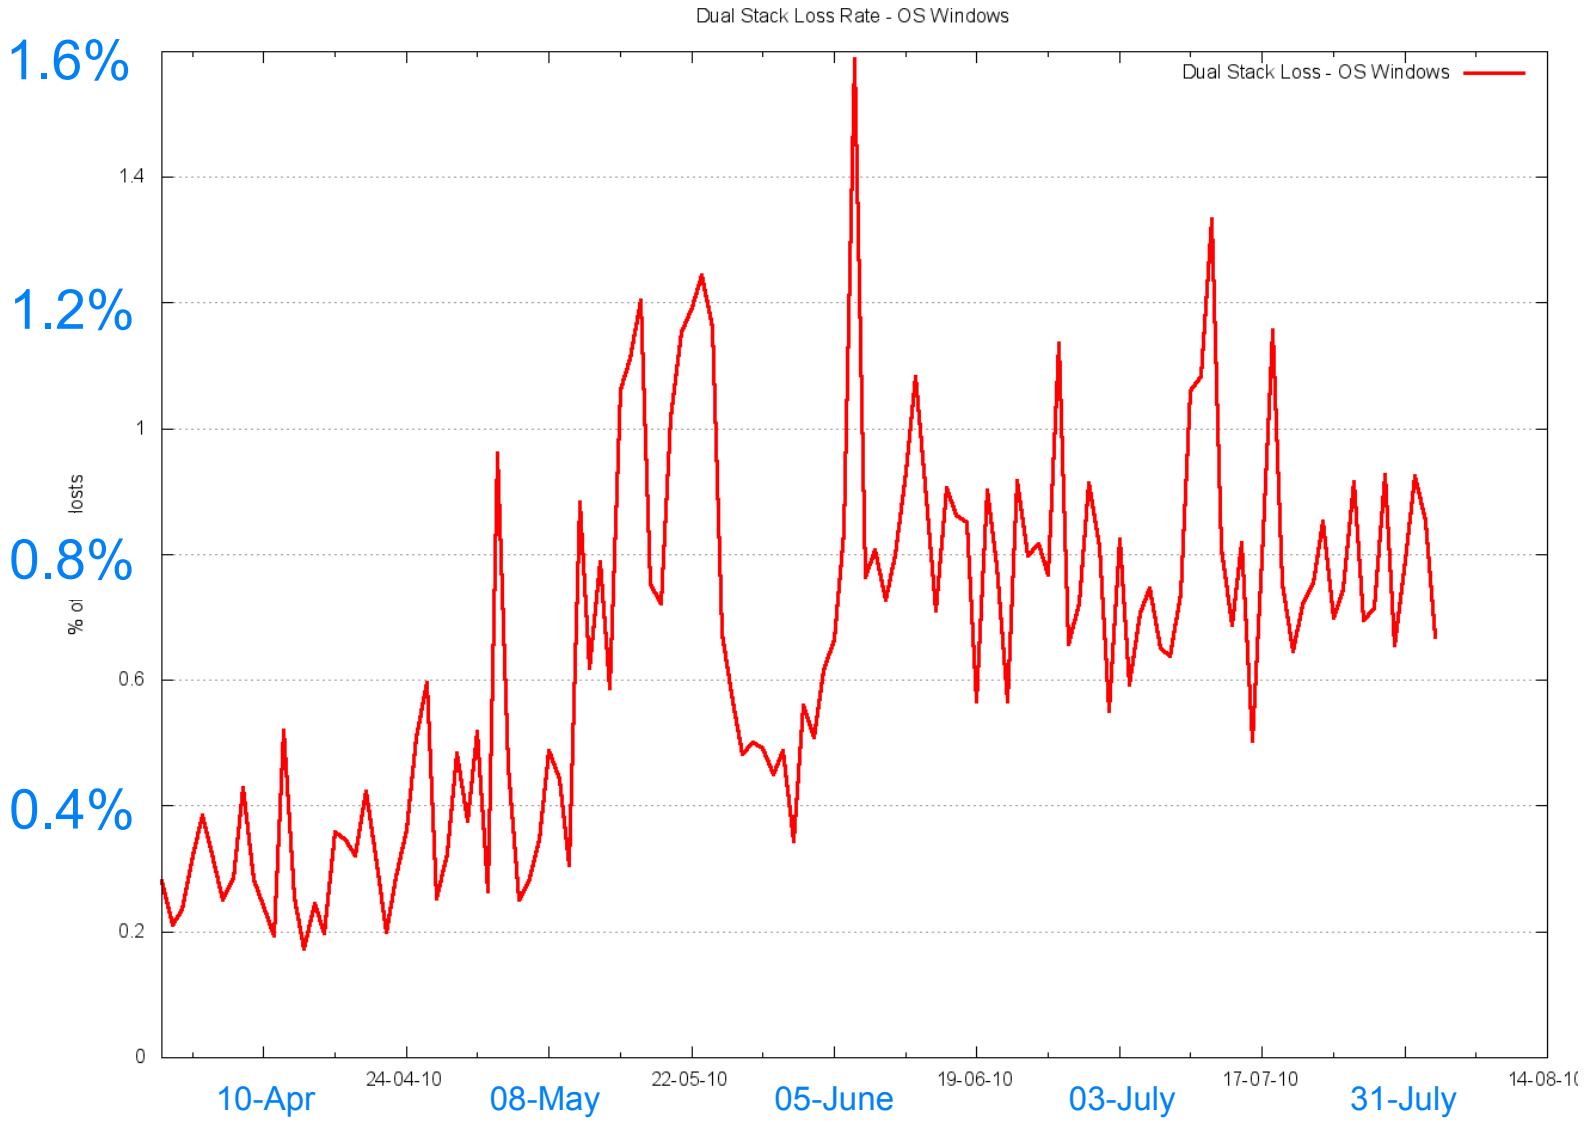

Dual Stack Failure

- One possible explanation of dual stack failure: IPv6 mis-configuration
  - The client can retrieve V4 objects
  - The client has a local V6 unicast address
  - When presented with a dual stack object the client will attempt to use V6 to retrieve the object
  - The V6 environment is misconfigured
  - The client times out on attempting to retrieve the V6 object, and does not attempt to revert to V4.

Dual Stack Failure

- Another possible explanation is user interruption
  - the script itself is loaded from a dual stack server, so the client was demonstrably able to perform a retrieval from a dual stack server
  - So the failure on loading the dual stack object could be due to user reset of the client application
- No clear picture as yet on why we are seeing this kind of client behaviour

Can you help?

- This data is collected by adding a small code fragment to your web site:

```
<script src="http://www.potaroo.net/linktest-js.php" type="text/javascript"></script>
```

- If you would like to assist us in this activity, please let us know:

research@apnic.net

Daily Reports

This is an on-going experiment which we will continue to operate throughout this transition

Daily Reports and the generated data sets are at:

http://www.potaroo.net/stats/1x1

Thank You!