Measuring IPv6 Day

Geoff Huston
APNIC

My brief for this session...

It would be great if you could consider to include following topics in your presentation:

What you observed on World IPv6 day:

- Statistics and measurements, and analysis (traffic, transition technology, performance comparison, DNS, OS, Browser, country/ regional break down etc.) on the day of W6D
- What worked?
- What did not work?
- What should the AP region do next?

  etc.

So what did we see on June 8 2011?

That we did not see on June 7 or June 9

Nothing!

Because:

- We are measuring the IPv6 behaviour of clients, not servers
- And, on the whole, clients did not change anything they were doing on World IPv6 Day

There is a story about measuring clients' capabilities and performance with IPv6.

But client behaviour is not a one day event story

It's a story about longer term behaviors and trends...

Some Dual Stack Questions

- How many clients are capable of IPv6 access?
- What forms of IPv6 access are they using?
- Is their experience over Dual Stack better or worse than IPv4?

An Approach to IPv6 Measurement

- We'd like a tool that will measure your clients' IPv6 behavior without having to add IPv6 to your website
  - Leverage cross-site URL fetches

- We'd like to integrate these measurements into existing tracking methods, and analytics framework
  - No new analysis tools needed

APNIC's IPv6 capability measurement system

http://labs.apnic.net

- Built on google 'analytics' method
  - Javascript, highly portable
  - Asynchronous, runs in the background
- Data integrated into Google Analytics reports
  - Graphs of 'events' to monitor IPv4, IPv6 and dual-stack
- Configurable by website manager
  - Sample or every connection, extra tests etc

Additional Measurements

We extended this technique into Flash, and created an anonymous banner ad

The IPv6 capability test is built into the Flash code

Banner Ad Fun

No clicks needed

(indeed we would prefer that clients did NOT click the ad, as it costs us more for a click!)

Impressions are really cheap

$25 per day buys around 25,000 impressions

Every impression carries the complete IPv6 test set

But many users are ad-intolerant

Users tend to browse away from pages containing the ad in a far shorter time interval

We see a higher number of aborted test runs with the ad

IPv6 capability, as seen by Google

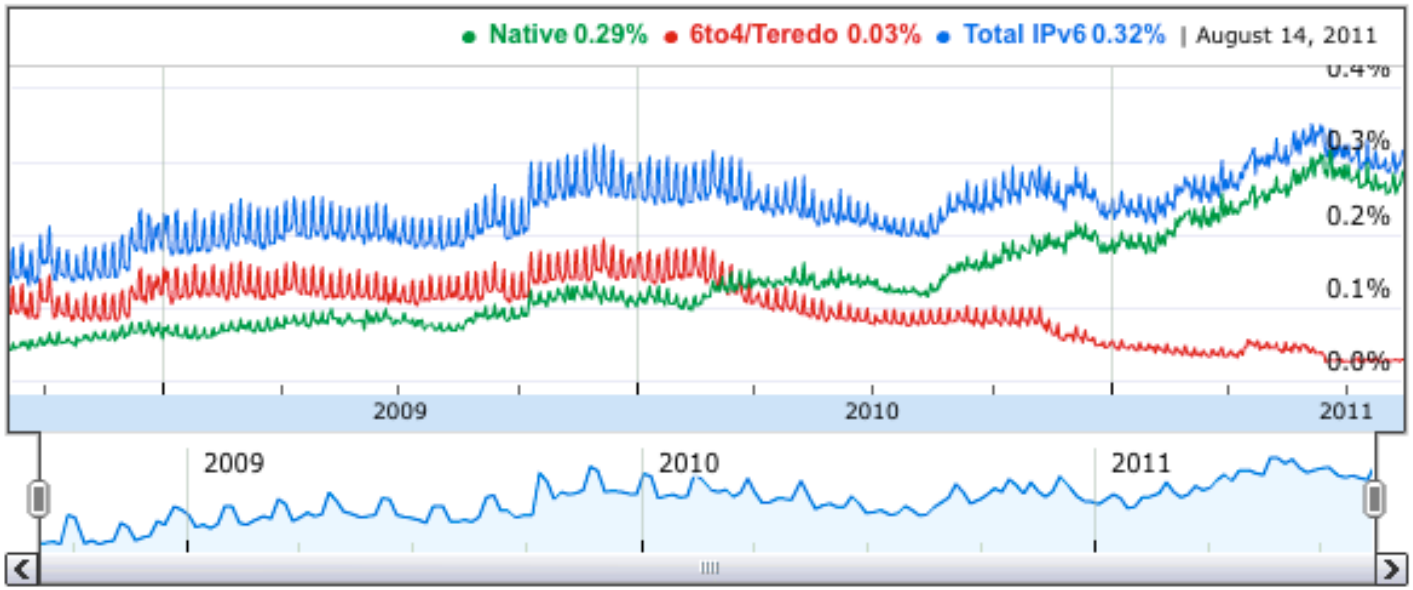

http://www.google.com/intl/en/ipv6/statistics/

IPv6 capability, as seen by APNIC

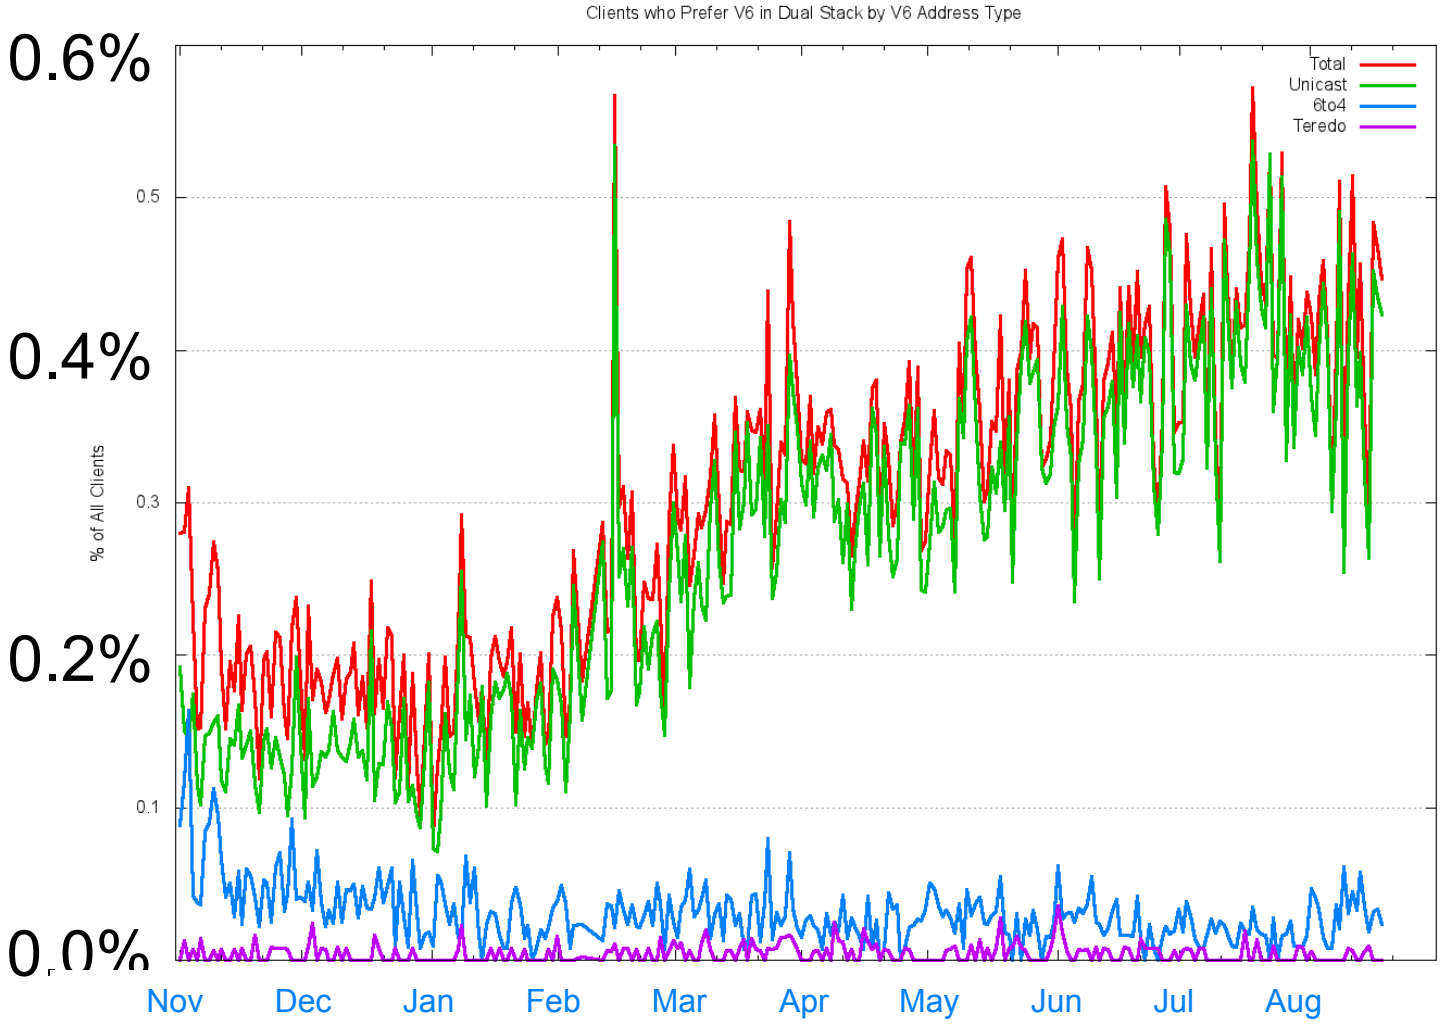

IPv6 capability, as seen by APNIC

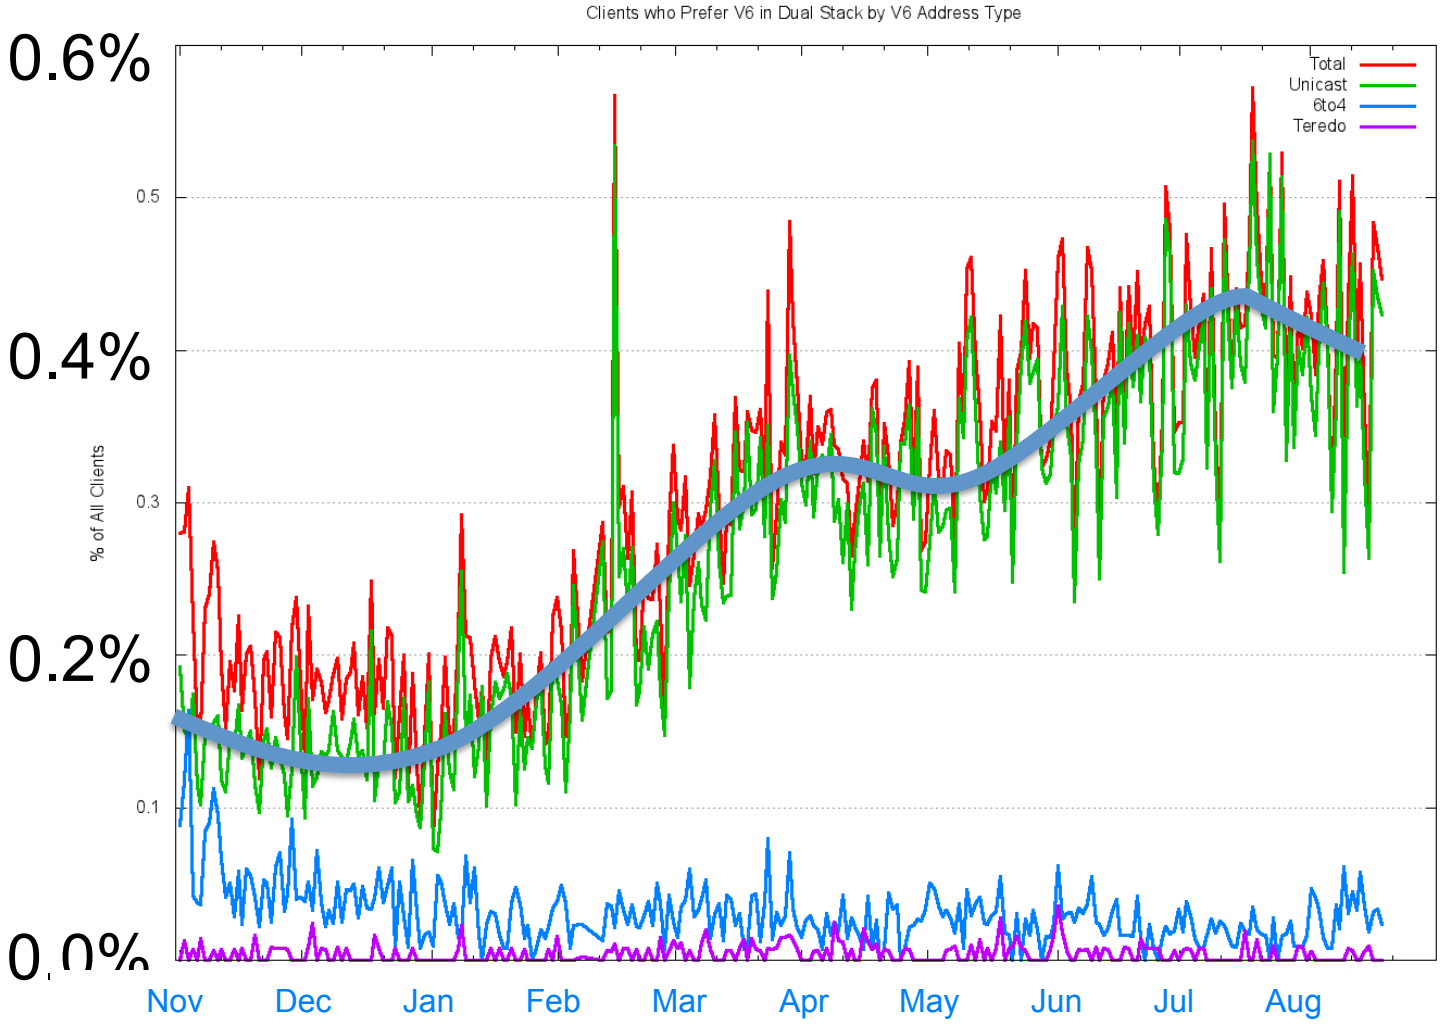

Is This All There Is?

- **0.3% – 0.4%** of clients is a very low number
  - And most of the IPv6 access we see here uses unicast IPv6
  - Where are all the 6to4 and Teredo auto-tunnels?
  - What is going on in the past few weeks with the drop in IPv6 access?
- Lets look harder by testing with an IPv6-only image

IPv6 Capable Clients

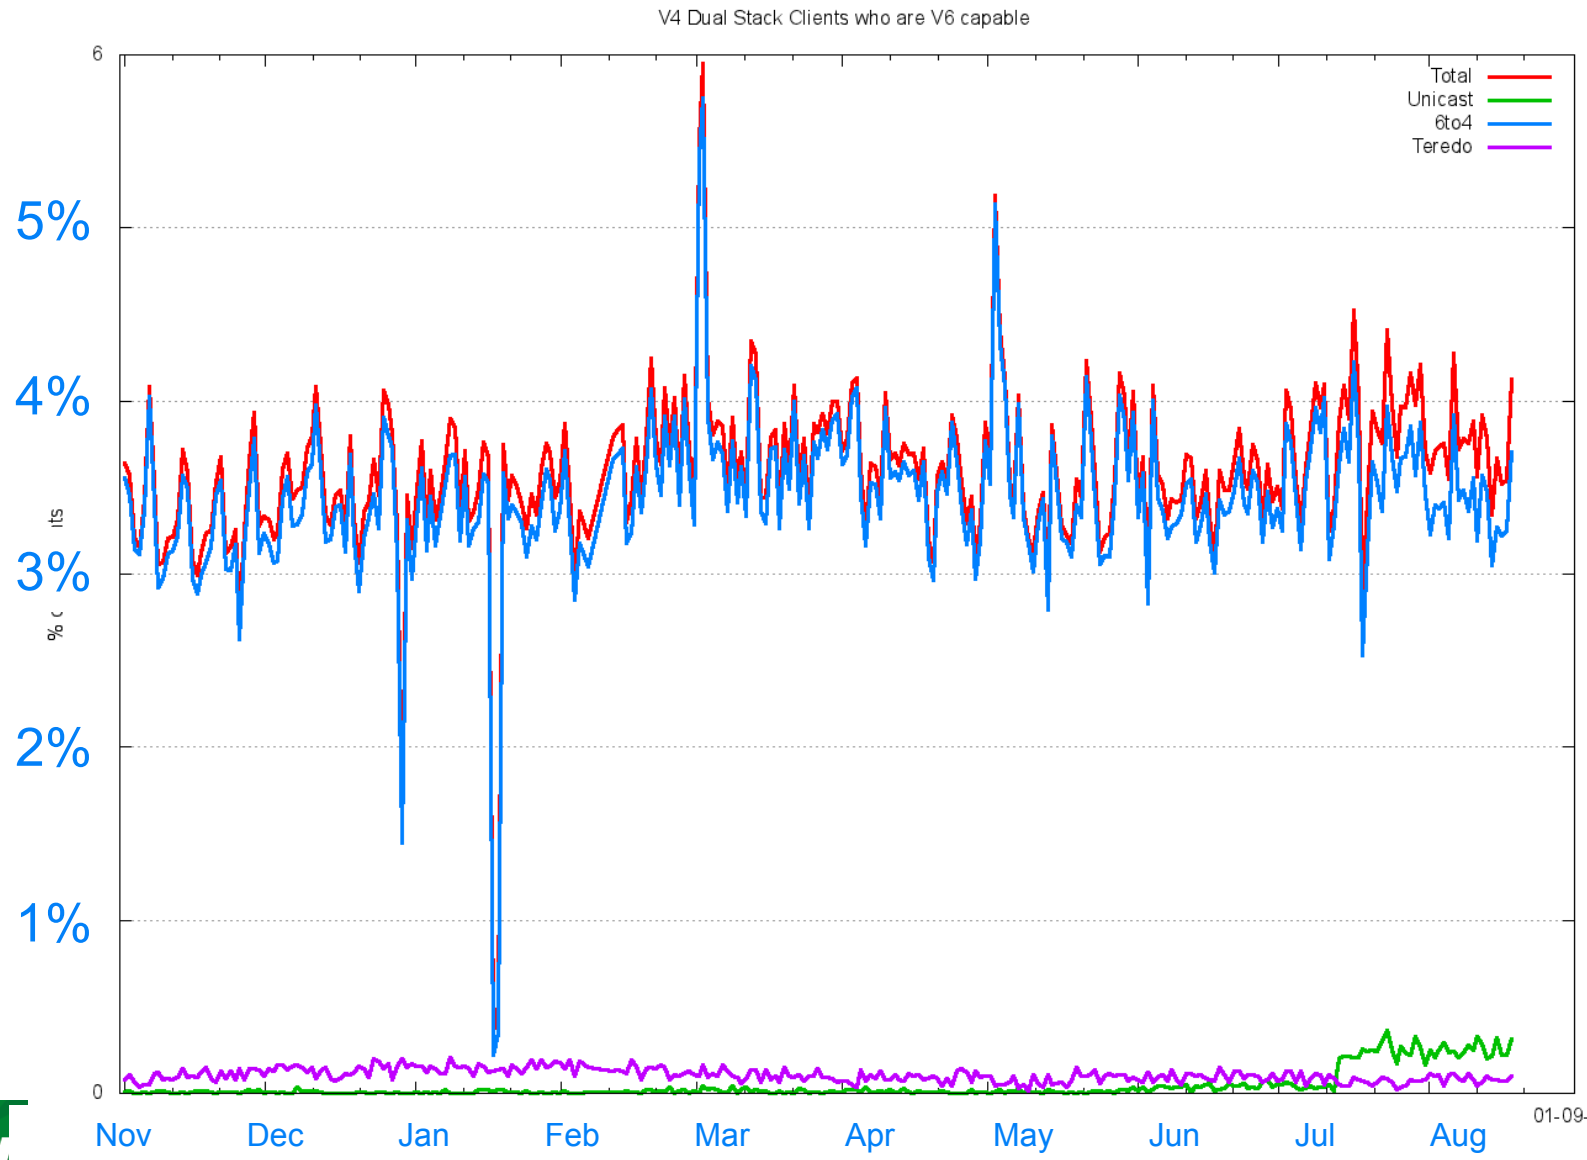

IPv6 Capable Clients

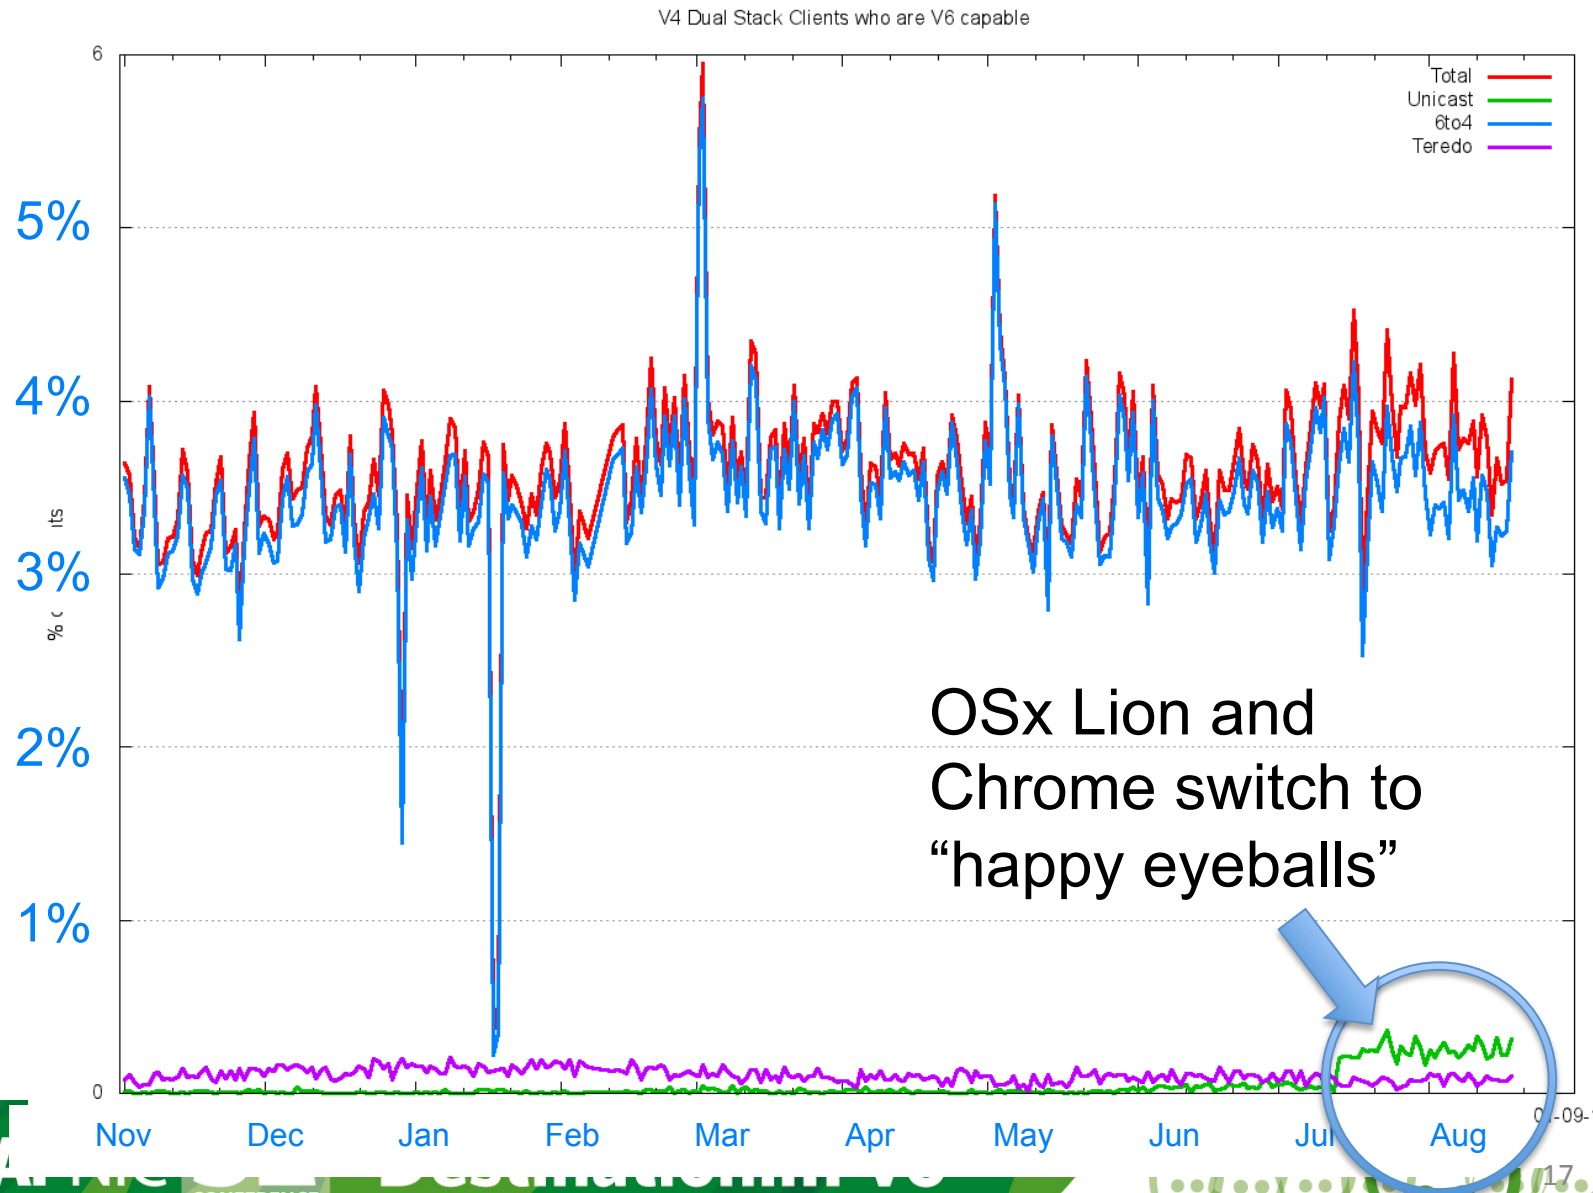

IPv6: “could” vs “will”

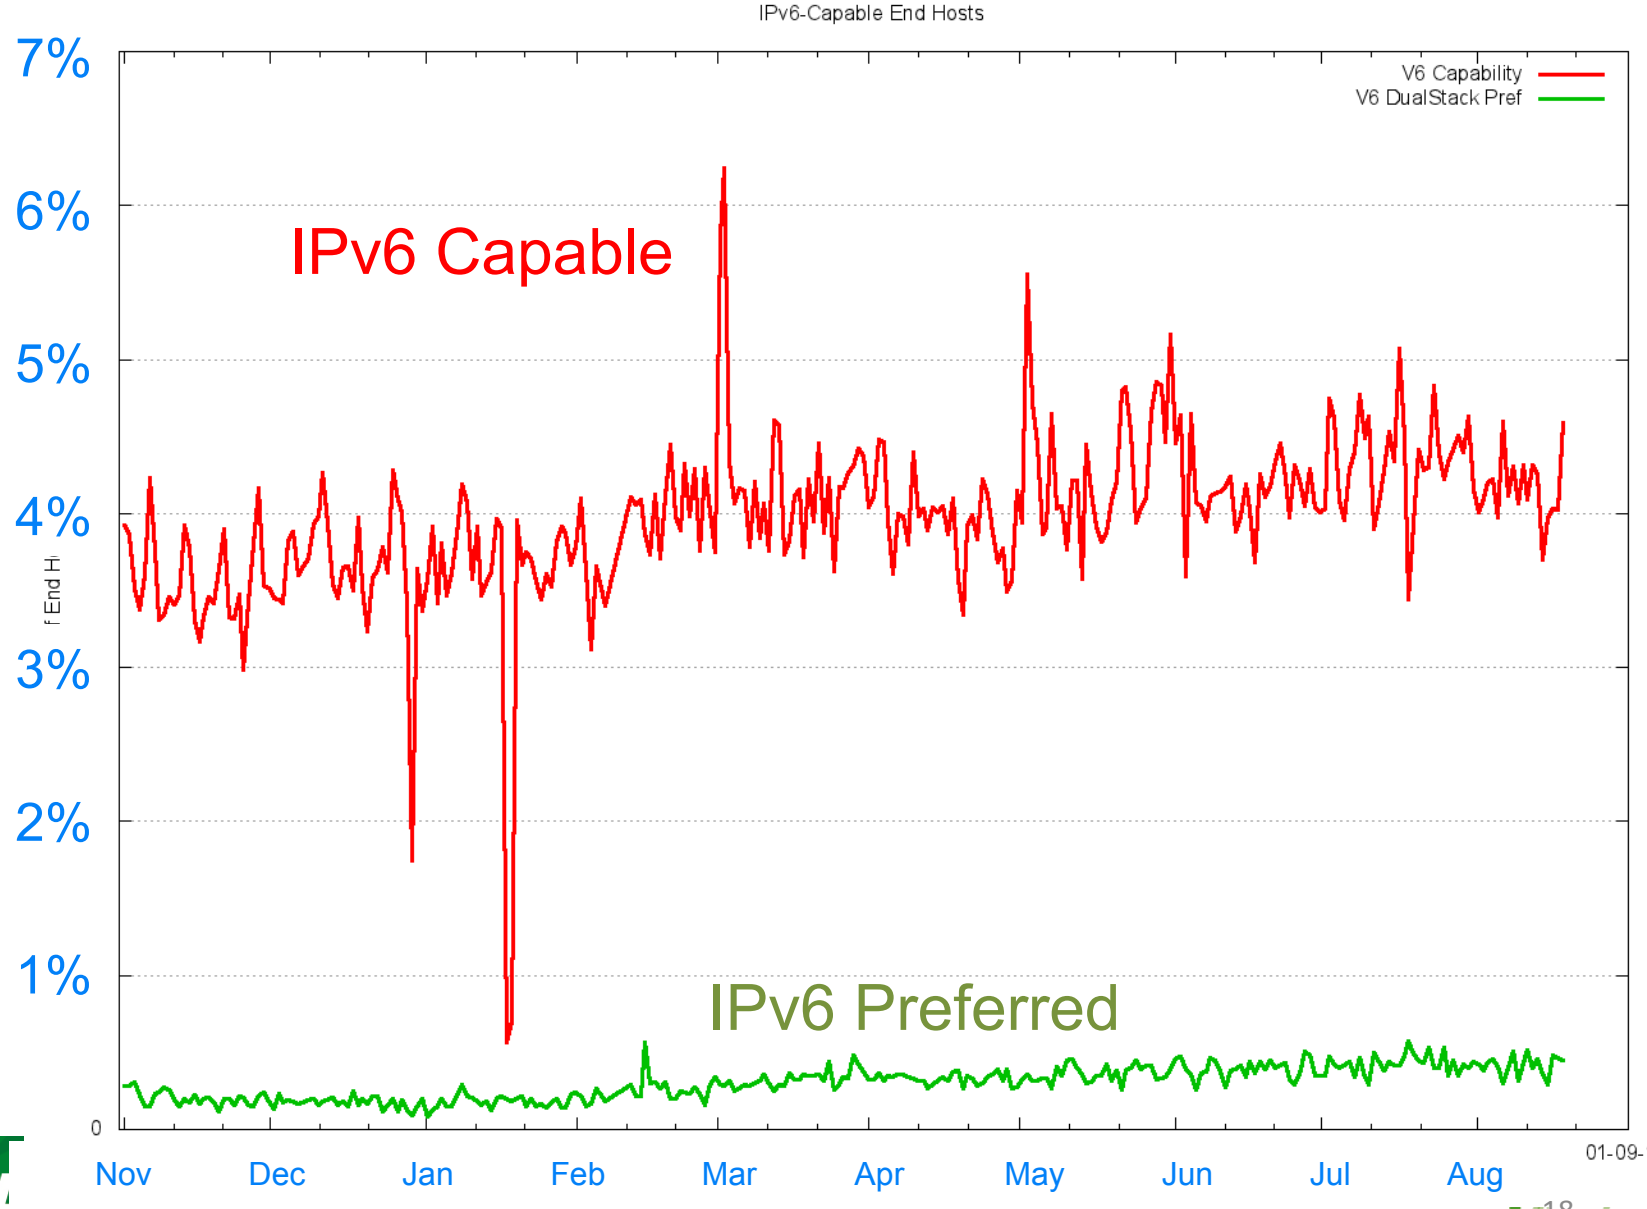

Is This All There Is?

- **3% - 4%** of clients is still a very low number
  - Most of the access in IPv6-only is via 6to4 auto-tunnelling
  - Where is Teredo?
- Lets look harder by testing with an image that does not require a DNS lookup:

  http://[2401:2000:6660::f003]/1x1.png

IPv6 Coerceable Clients

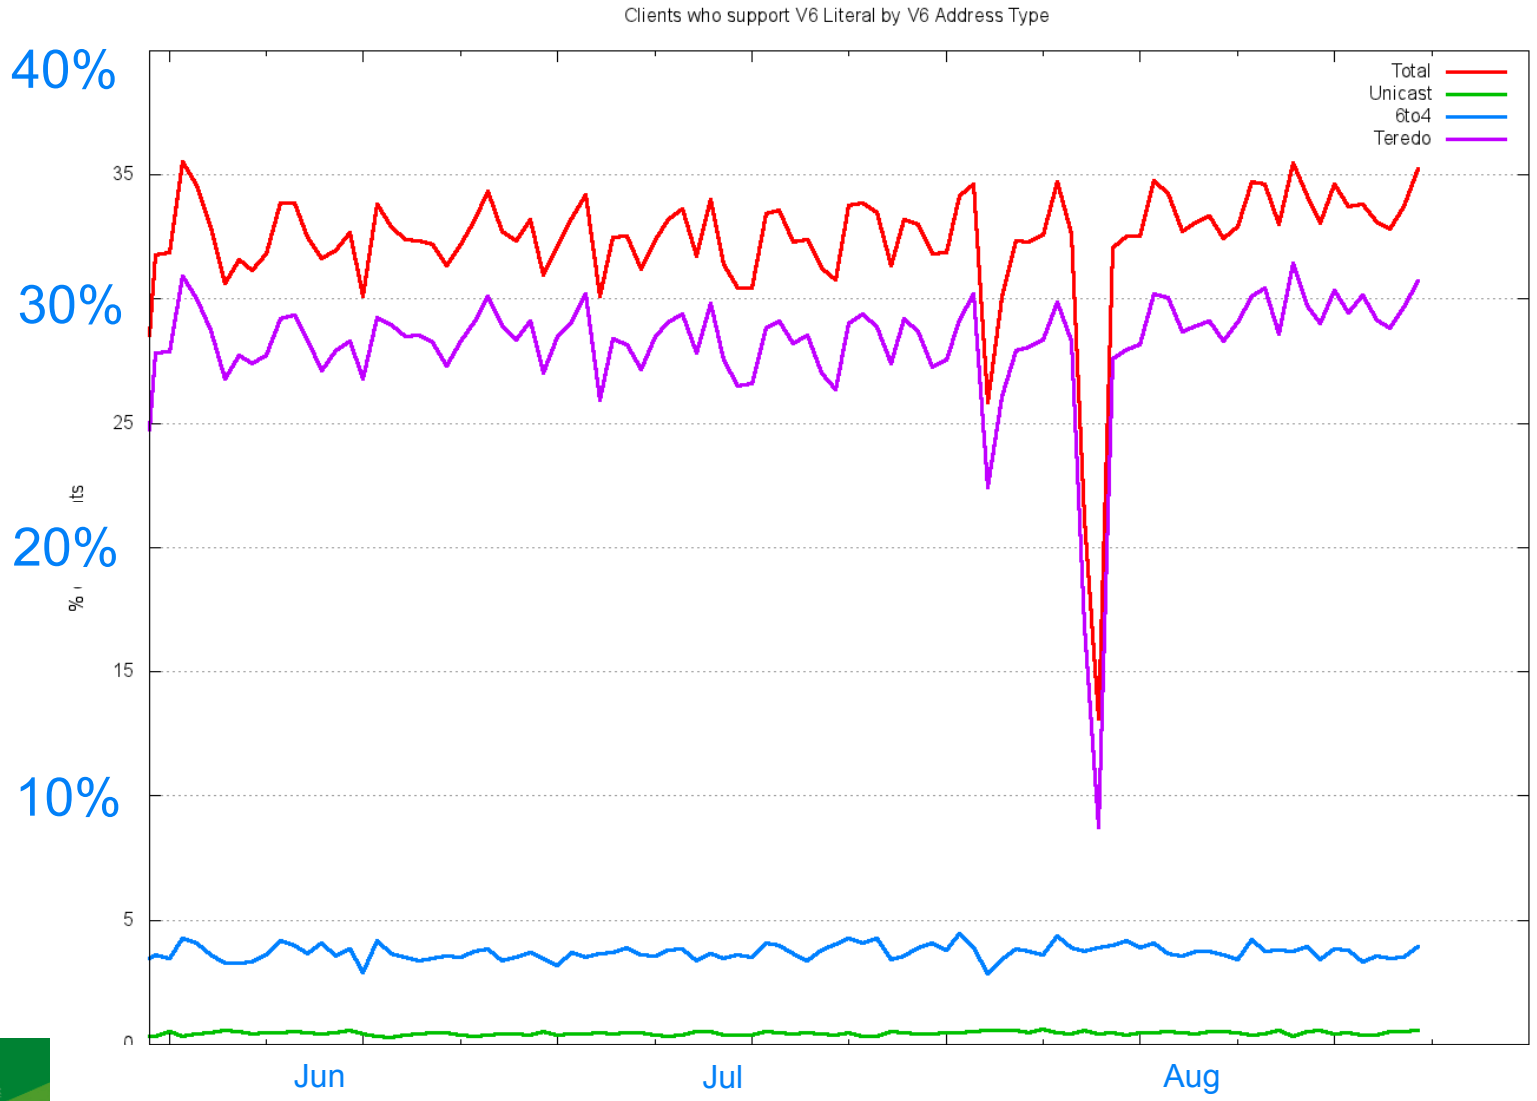

IPv6 Client Capabilities

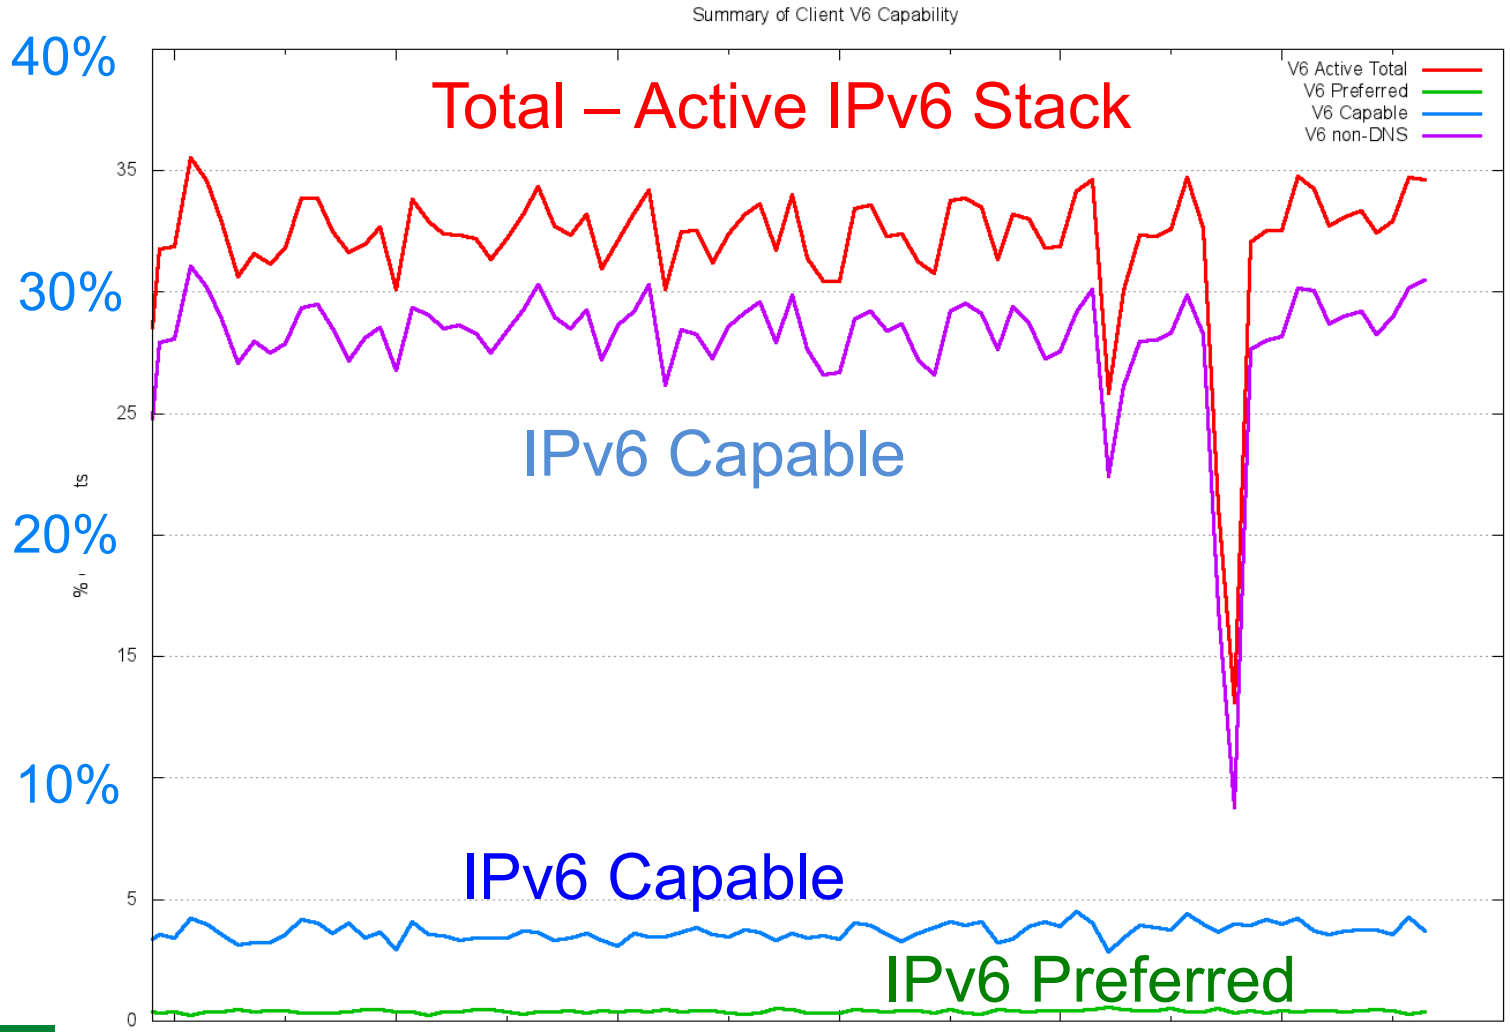

How Much IPv6 is Out There?

- Around **0.4%** of the Internet's clients can and will use IPv6 in a Dual Stack scenario
  - These clients are generally using a "native" IPv6 service
- Around **4%** of the Internet's clients can use IPv6 in an IPv6-only scenario
  - The additional clients are generally using 6to4 auto-tunnelling
- Around **28%** of the Internet's clients are equipped with IPv6 capability that can be exposed
  - The additional clients are using Teredo auto-tunnelling

Failure Observations

Dual Stack Loss Rate

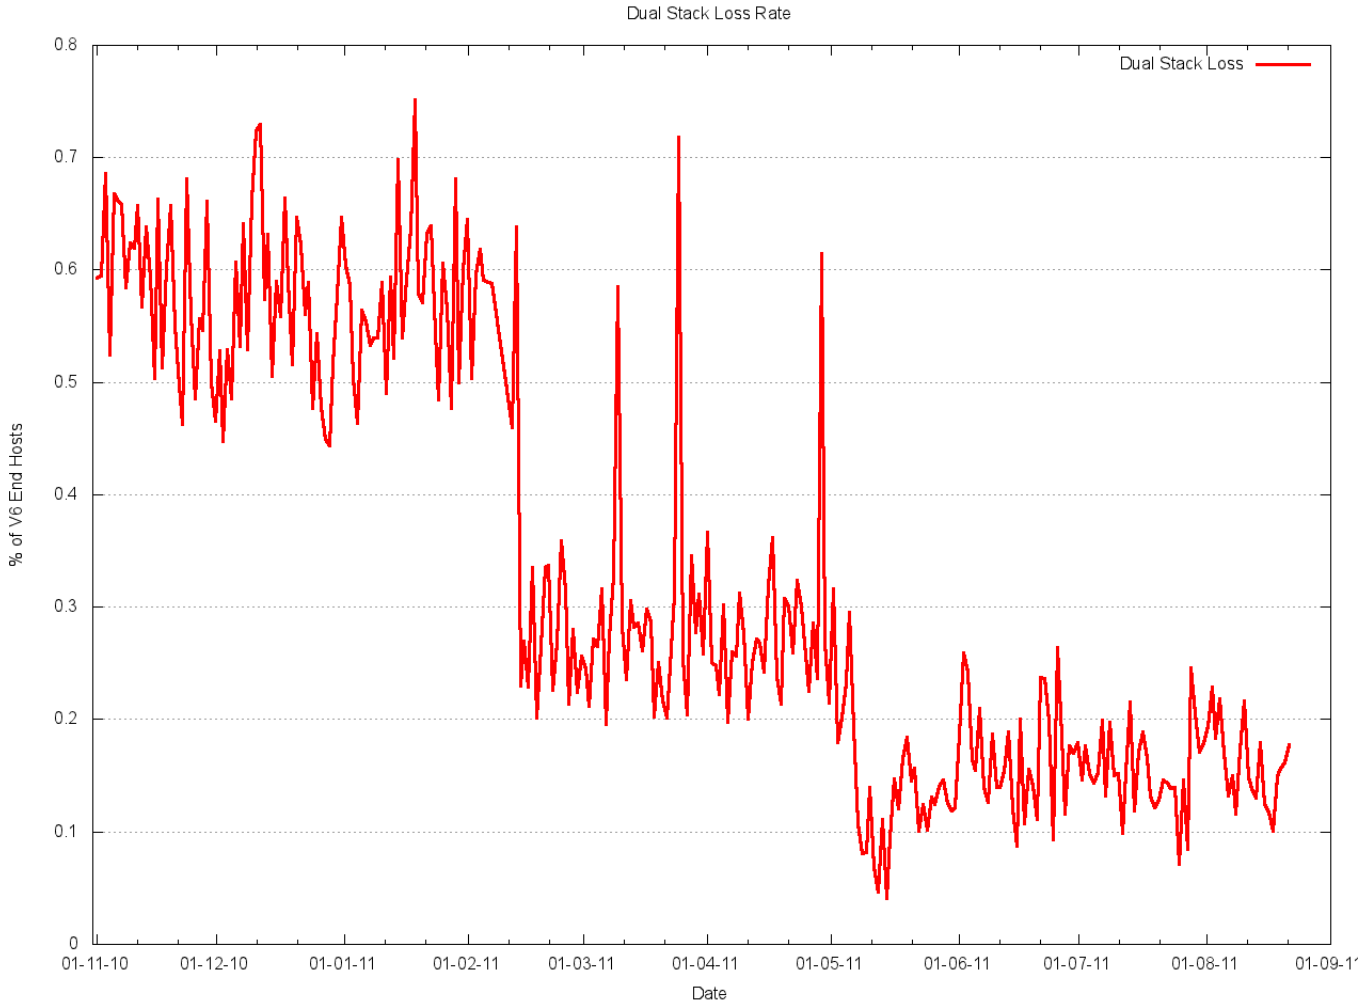

Dual Stack Loss Rate

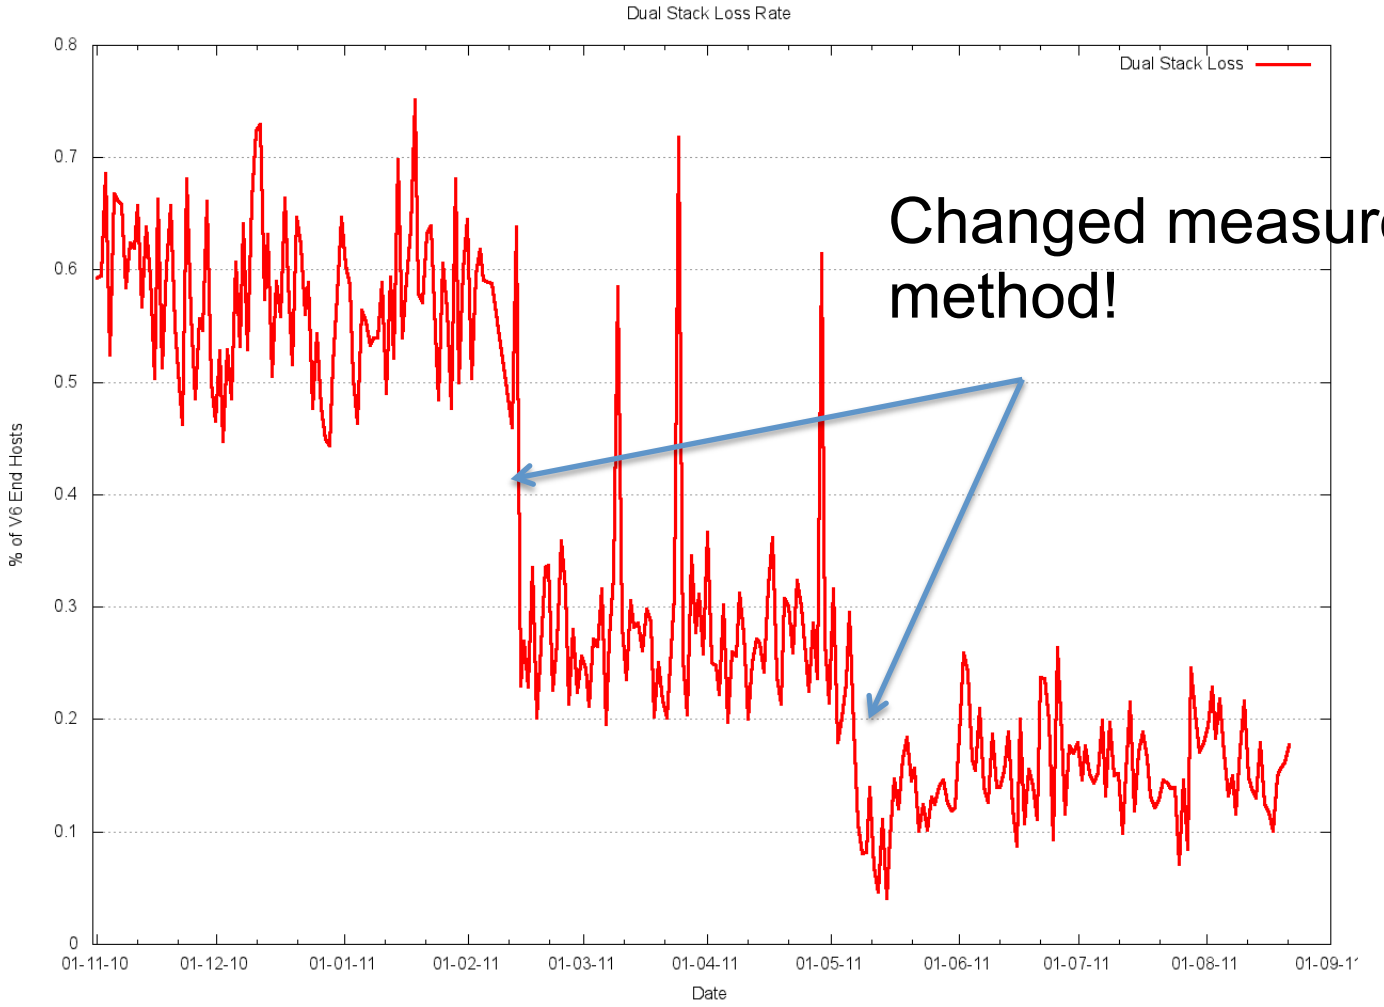

Changed measurement method!

Dual Stack Loss

- 15 in 10,000 clients are unable to fetch a web URL if presented with a dual-stack DNS name

  Likely from older (Windows XP?) hosts

- This is not very reliable data – other measurement exercises indicate a dual stack failure rate of around 3 per 10,000 clients

Connection Failure

To attempt to look more precisely for **some** instances of connection failure, lets looking for connections that fail after the initial TCP SYN

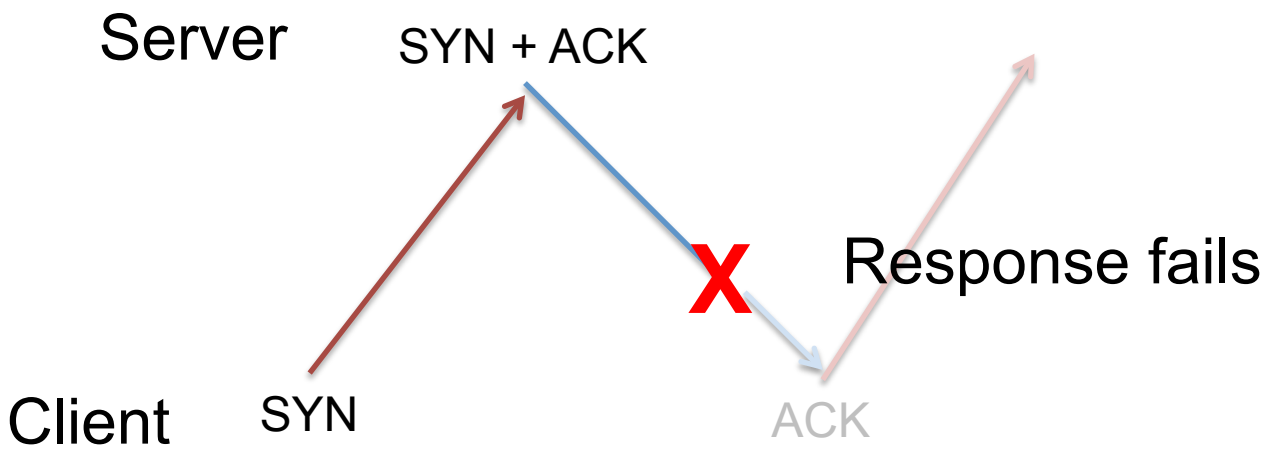

Note that this approach does not detect failure of the initial SYN packet, so the results are a lower bound of total connection failure rates

Connection Failure

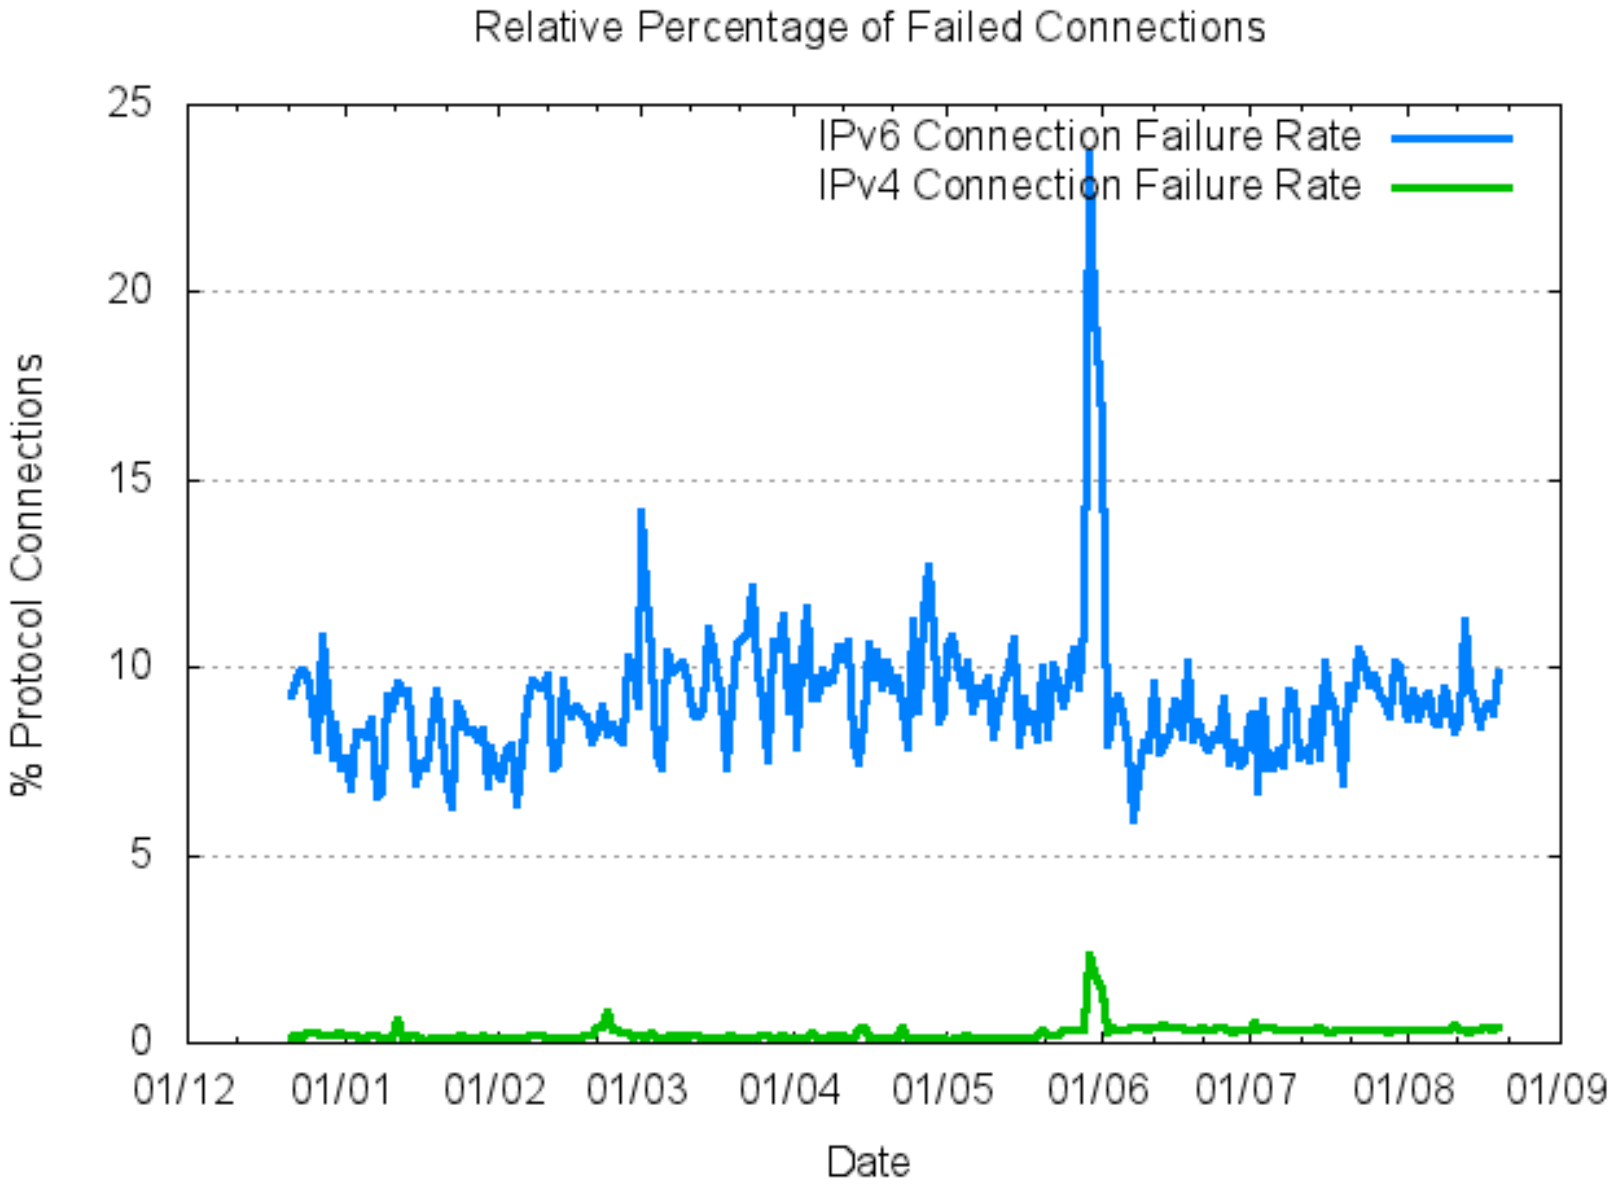

IPv6 Connection Failure

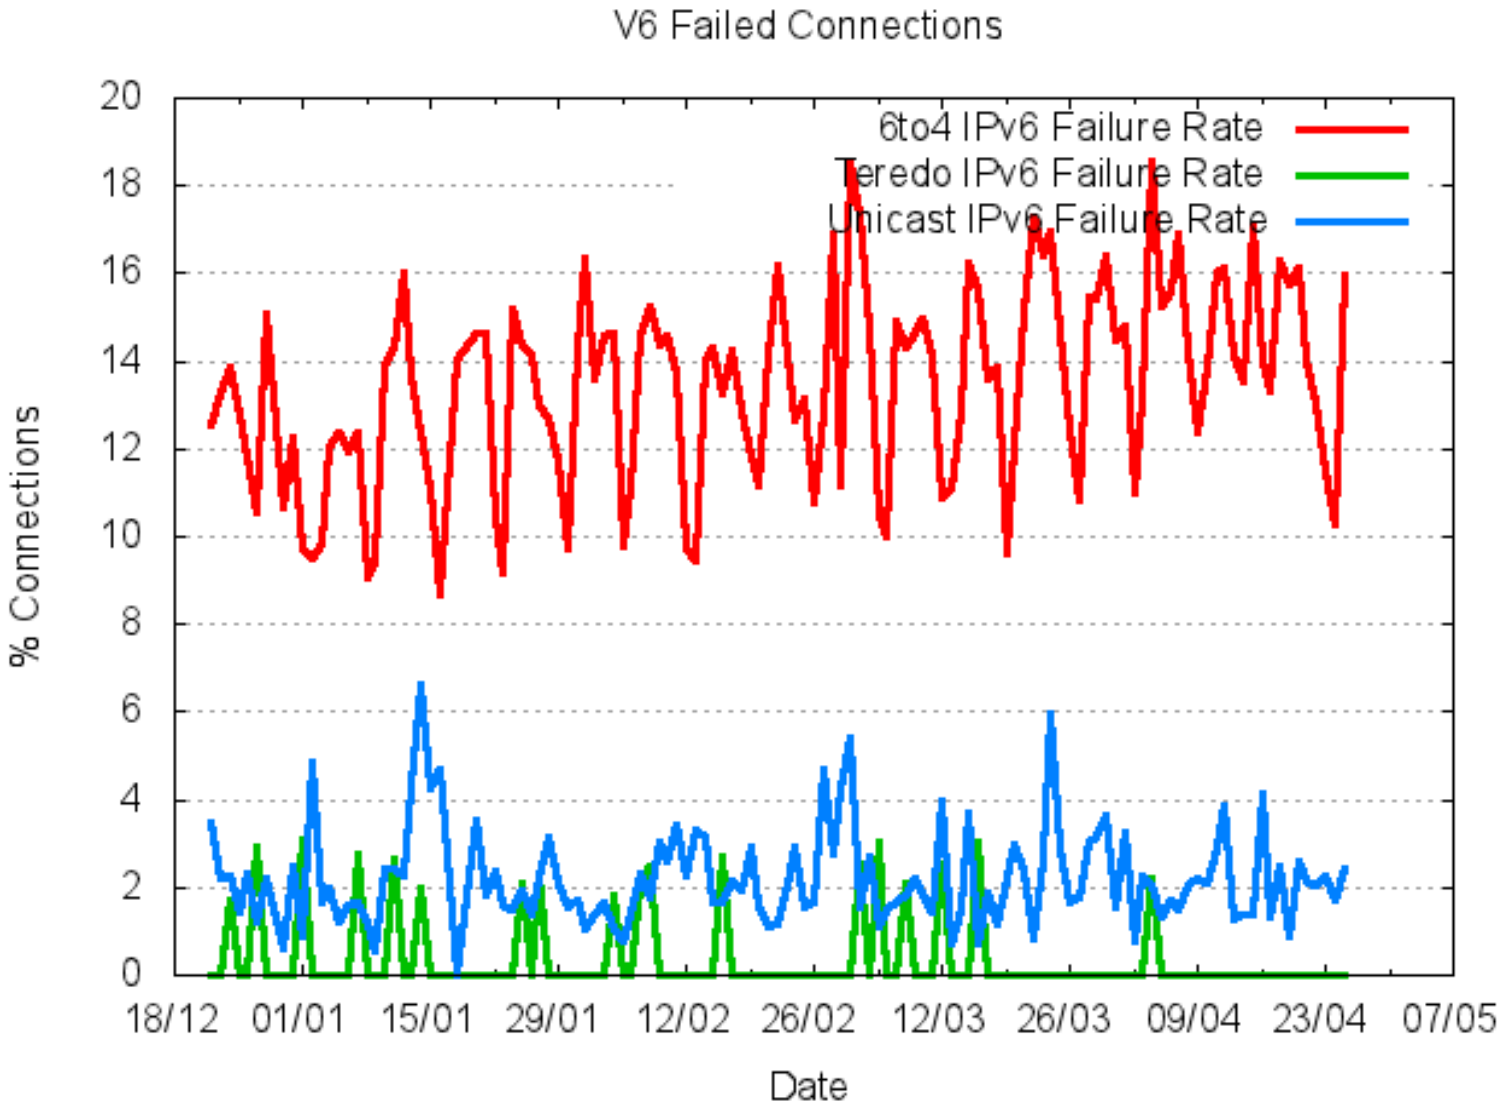

Is Teredo really THAT good?

Teredo Connection Failure

Teredo uses an initial ICMPv6 exchange to assist in the Teredo Server / Relay state setup

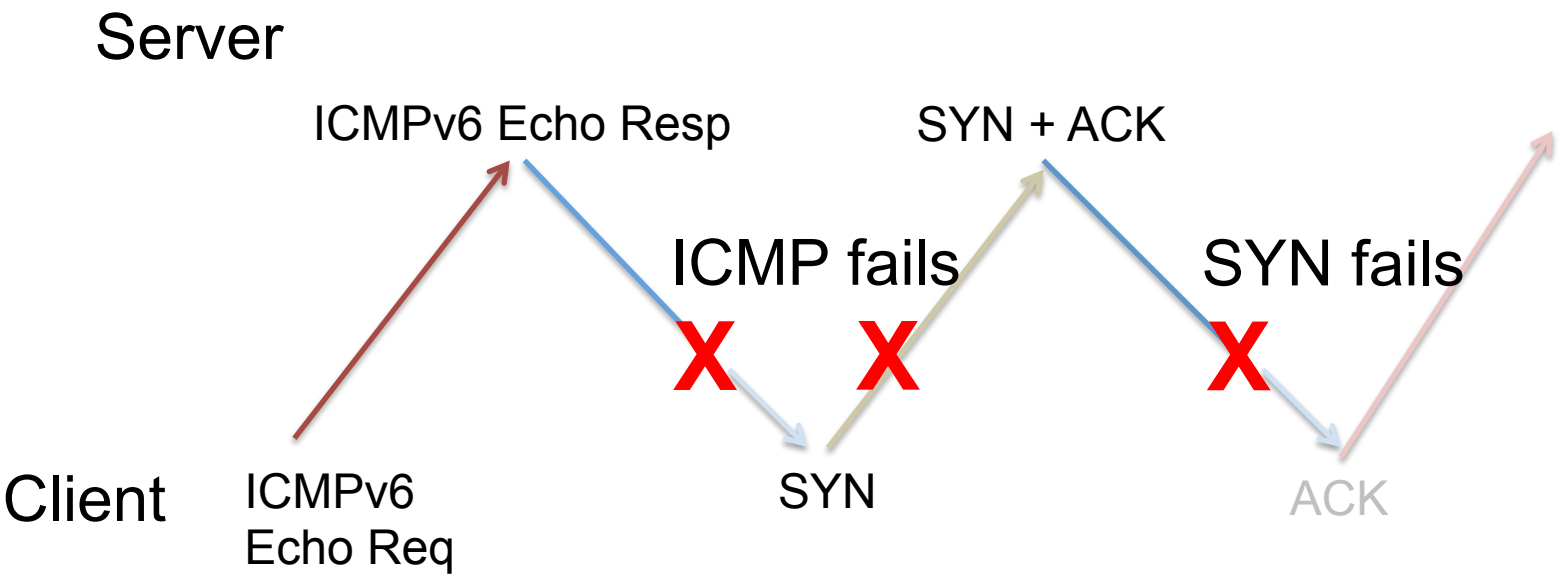

Note that this approach does not detect failure of the initial ICMPv6 echo request , so the results are a lower bound of total connection failure rates

IPv6 Connection Failure

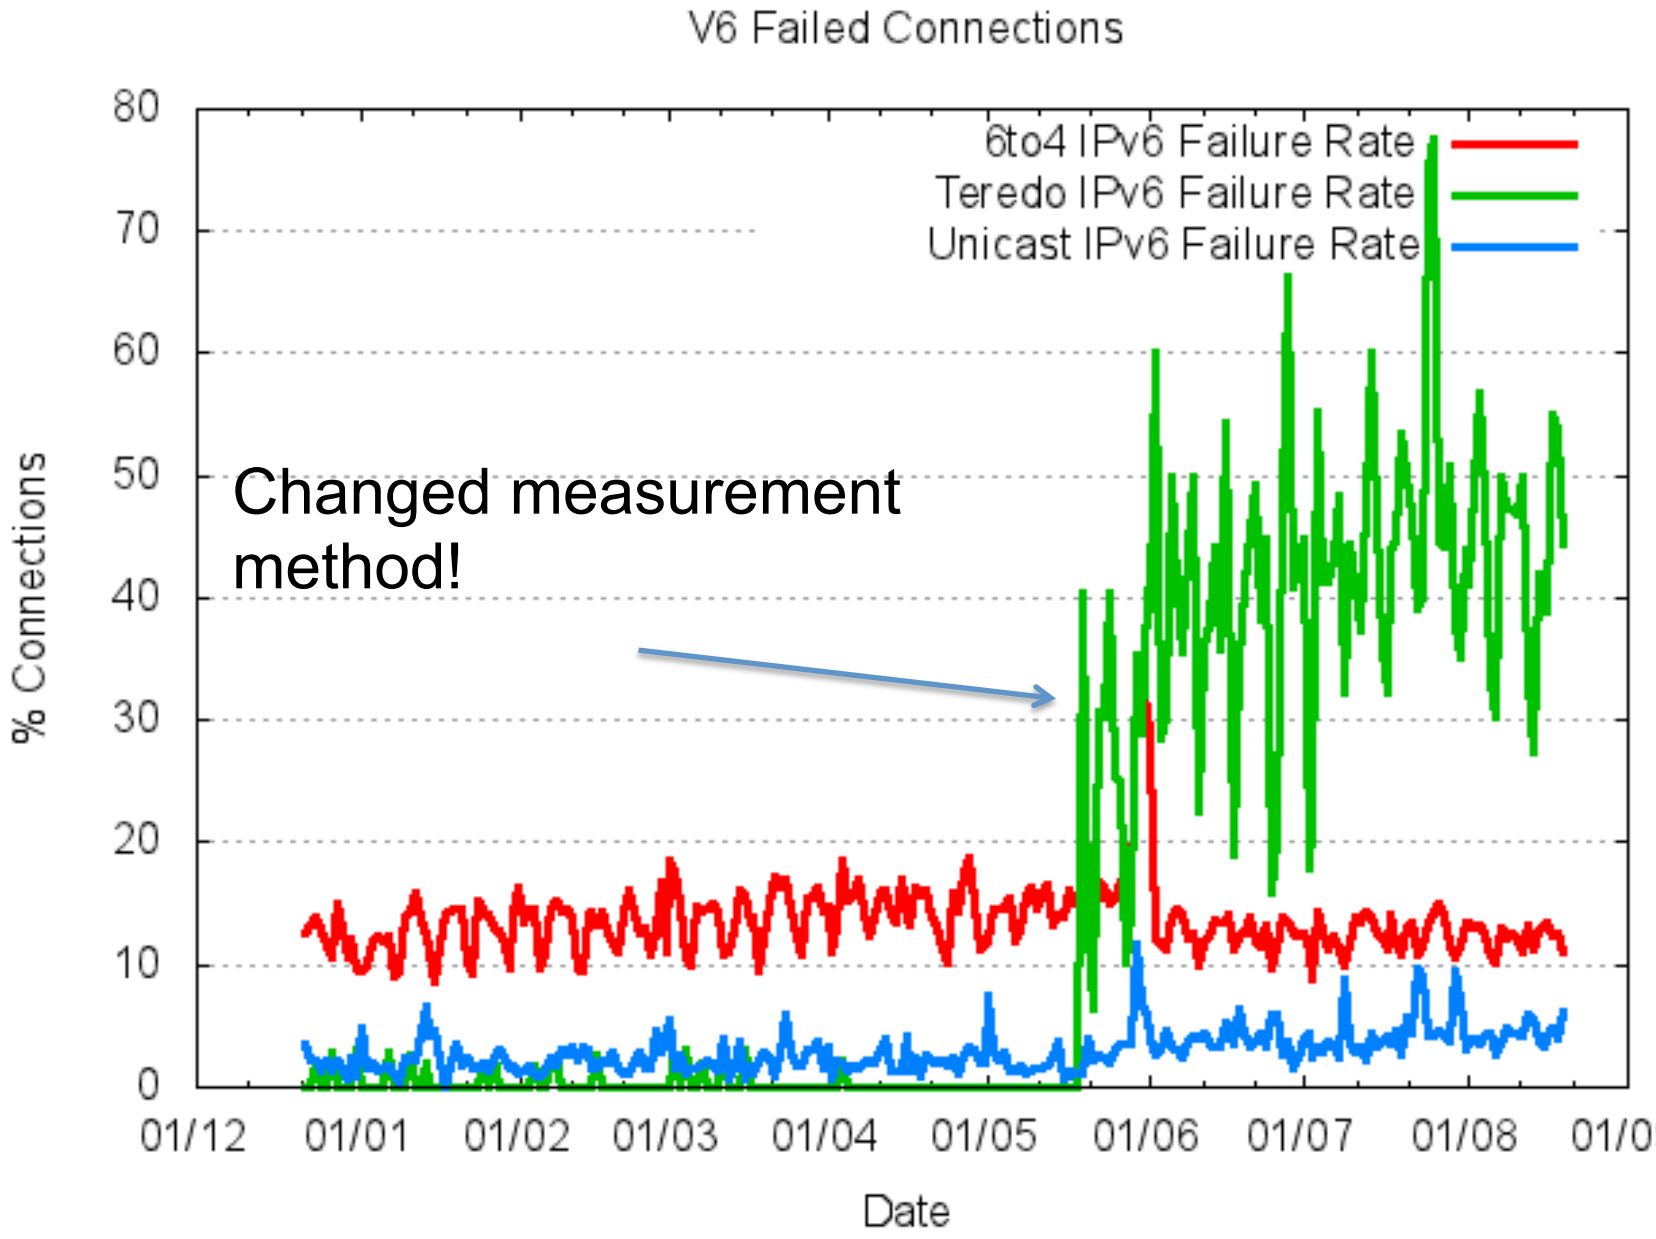

IPv6 Connection Failure

- Some **2%-5%** of **IPv6 unicast** connections fail!
  - This rate is better than IPv6 auto-tunnels, but is still 20x the rate of IPv4 connection failure
- Some **12% - 15%** of **6to4** connections fail!
  - This is a very high failure rate!
  - The failure is most likely a protocol 41 filter close to the client that prevents incoming 6to4 packets reaching the client
- Some **45%** of **Teredo** connections fail!
  - This is an amazingly high failure rate!
  - Is STUN just broken as a NAT traversal technology? Or …?

Conclusions

What can we say about the performance and robustness of a Dual Stack content service as a result of these observations?

For an Online Service…

Converting a service to operate as a Dual Stack service is a viable option in today's environment

And the latest Chrome and OS Lion dual stack client software makes this mode work as well as, or possibly faster, than IPv4!

But:

- a very small fraction (~0.05%) of existing clients will experience a much slower service through timeout failover from V6 to V4
- a very very small fraction (~0.03%) of existing clients will fail to connect to the dual stack service at all

What about Dual Stack Transition?

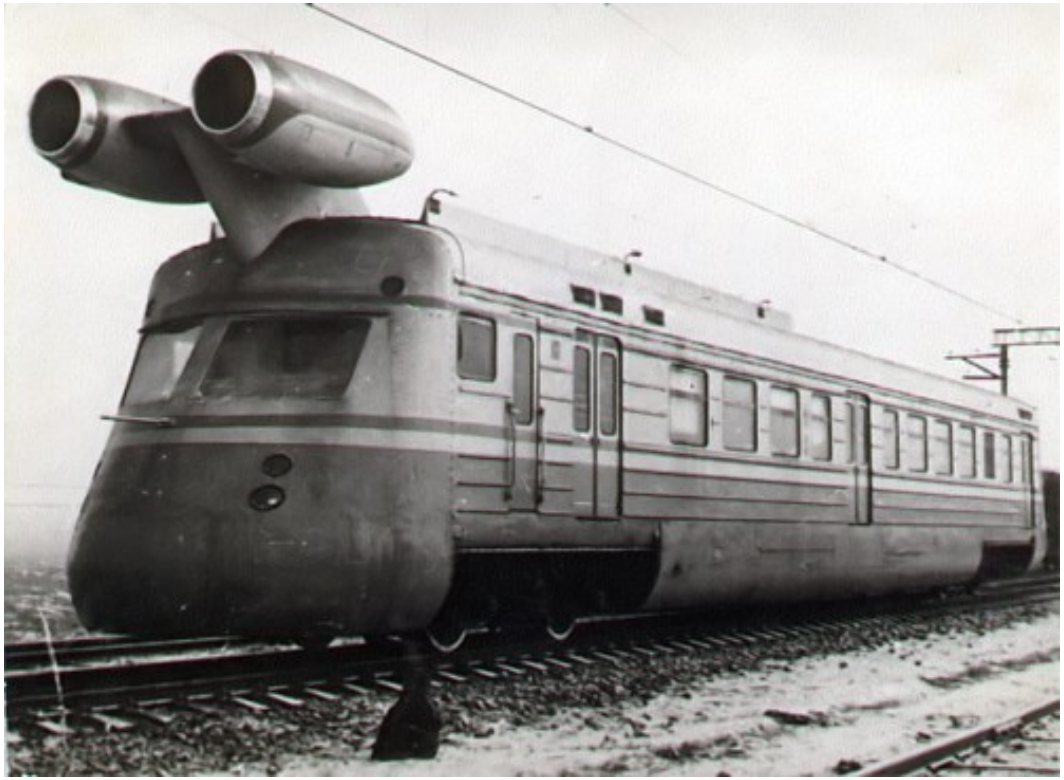

What about Dual Stack Transition?

End-host auto-tunnelling is not a solution!

- Auto-tunnelling appears to encounter many more performance and reliability problems than it solves in terms of IPv6 connectivity
- Auto-tunnelling is **not** proving to be a useful mainstream transition tool for IPv6

What about Dual Stack Transition?

If we want this transition to operate in a manner where IPv6 operates at least as well as IPv4 then end hosts really need to be connected to a IPv6 Unicast service delivered from their service provider

Thank You