

Measuring IPv6

Geoff Huston

APNIC R&D



Some IPv6 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

Insert an IPv6 “test” into a web page

- Whenever the client visits the web page the client will execute the “test”
- The test consists of a number of 1x1 gif element fetches
 - Dual Stack
 - IPv4 only
 - IPv6 only
 - Plus others....

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

But...

Measuring the IPv6 capabilities from a small number of web sites is not necessarily representative of the entire Internet (unless you are Google!)

So can we expand the measurement system to look at a broader sample of everyone?

The Power of Advertising!

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

A banner advertisement with a black background. The text "Are You IPv6 Ready?" is displayed in a stylized, multi-colored font. "Are" is pink, "You" is blue, "IPv6" is yellow, and "Ready?" is purple.

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!)

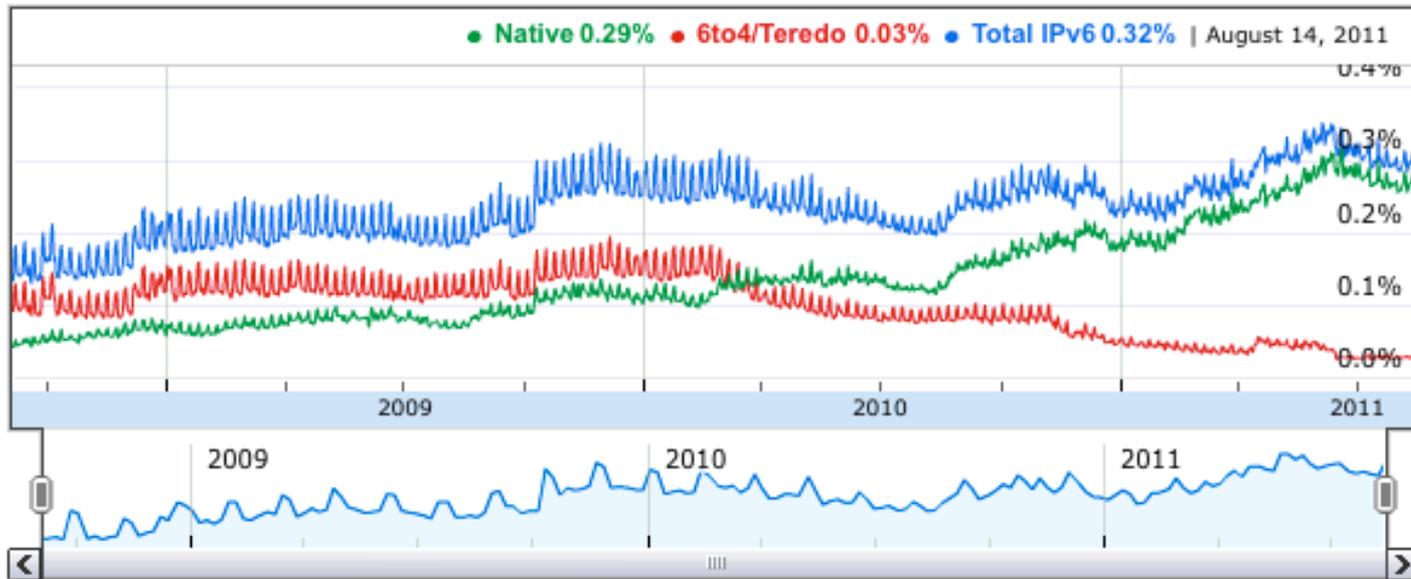
Impressions are really cheap

\$25 per day buys around 25,000 impressions

Every impression carries the complete IPv6 test set

And we get impressions from all over the Internet

IPv6 capability, as seen by Google

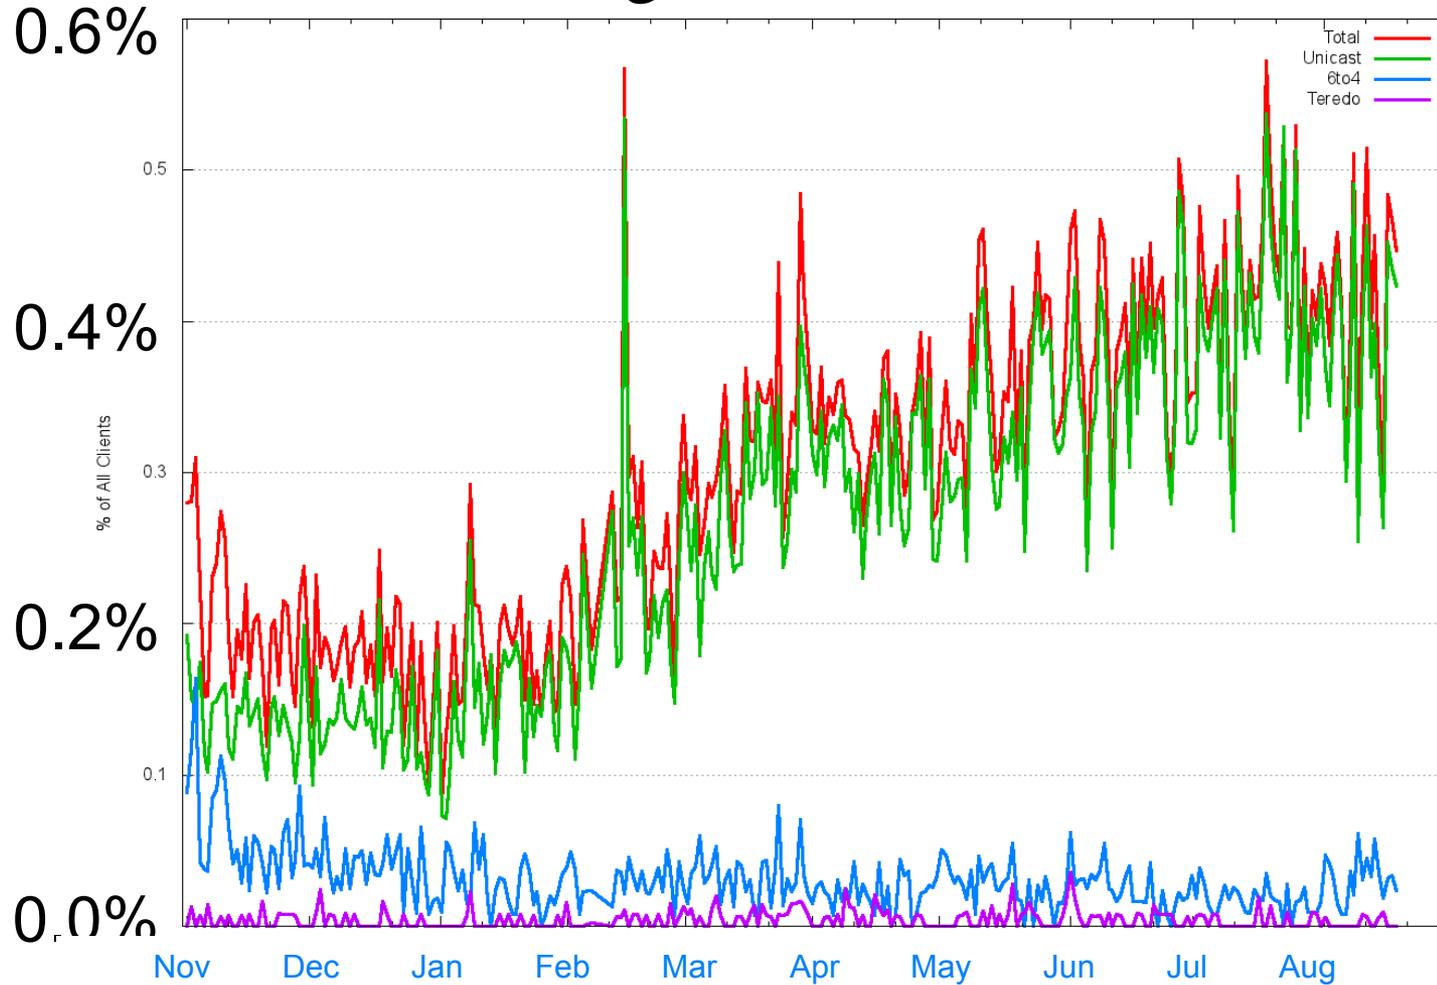


©2011 Google

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

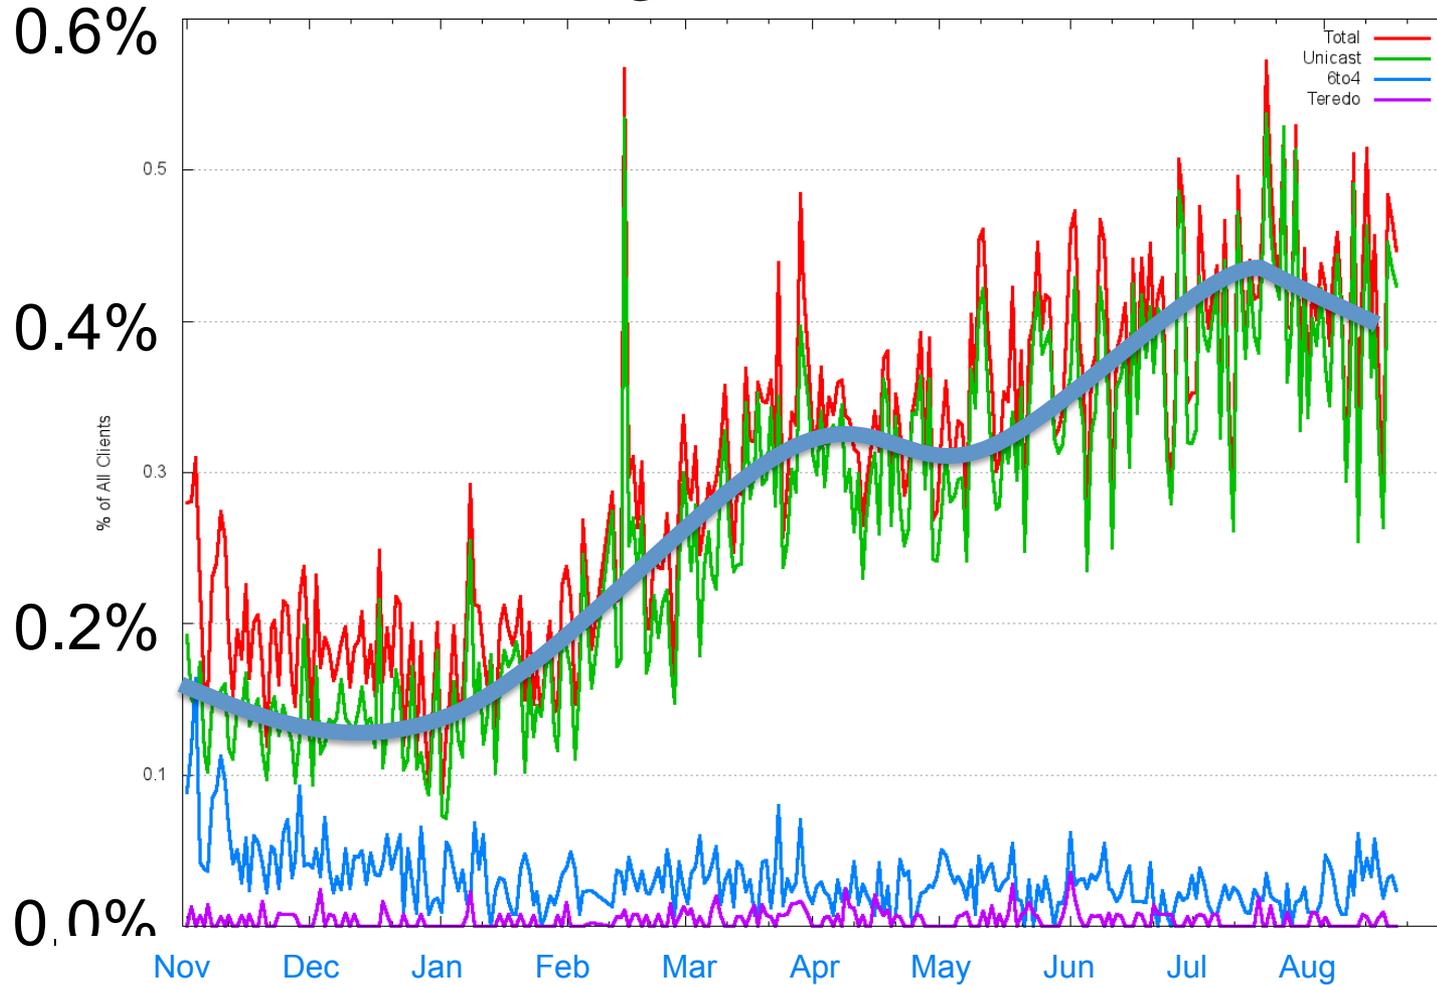
IPv6 capability, as seen by APNIC

Client no Prefer V6 in Dual Stack by V6 Address Type



IPv6 capability, as seen by APNIC

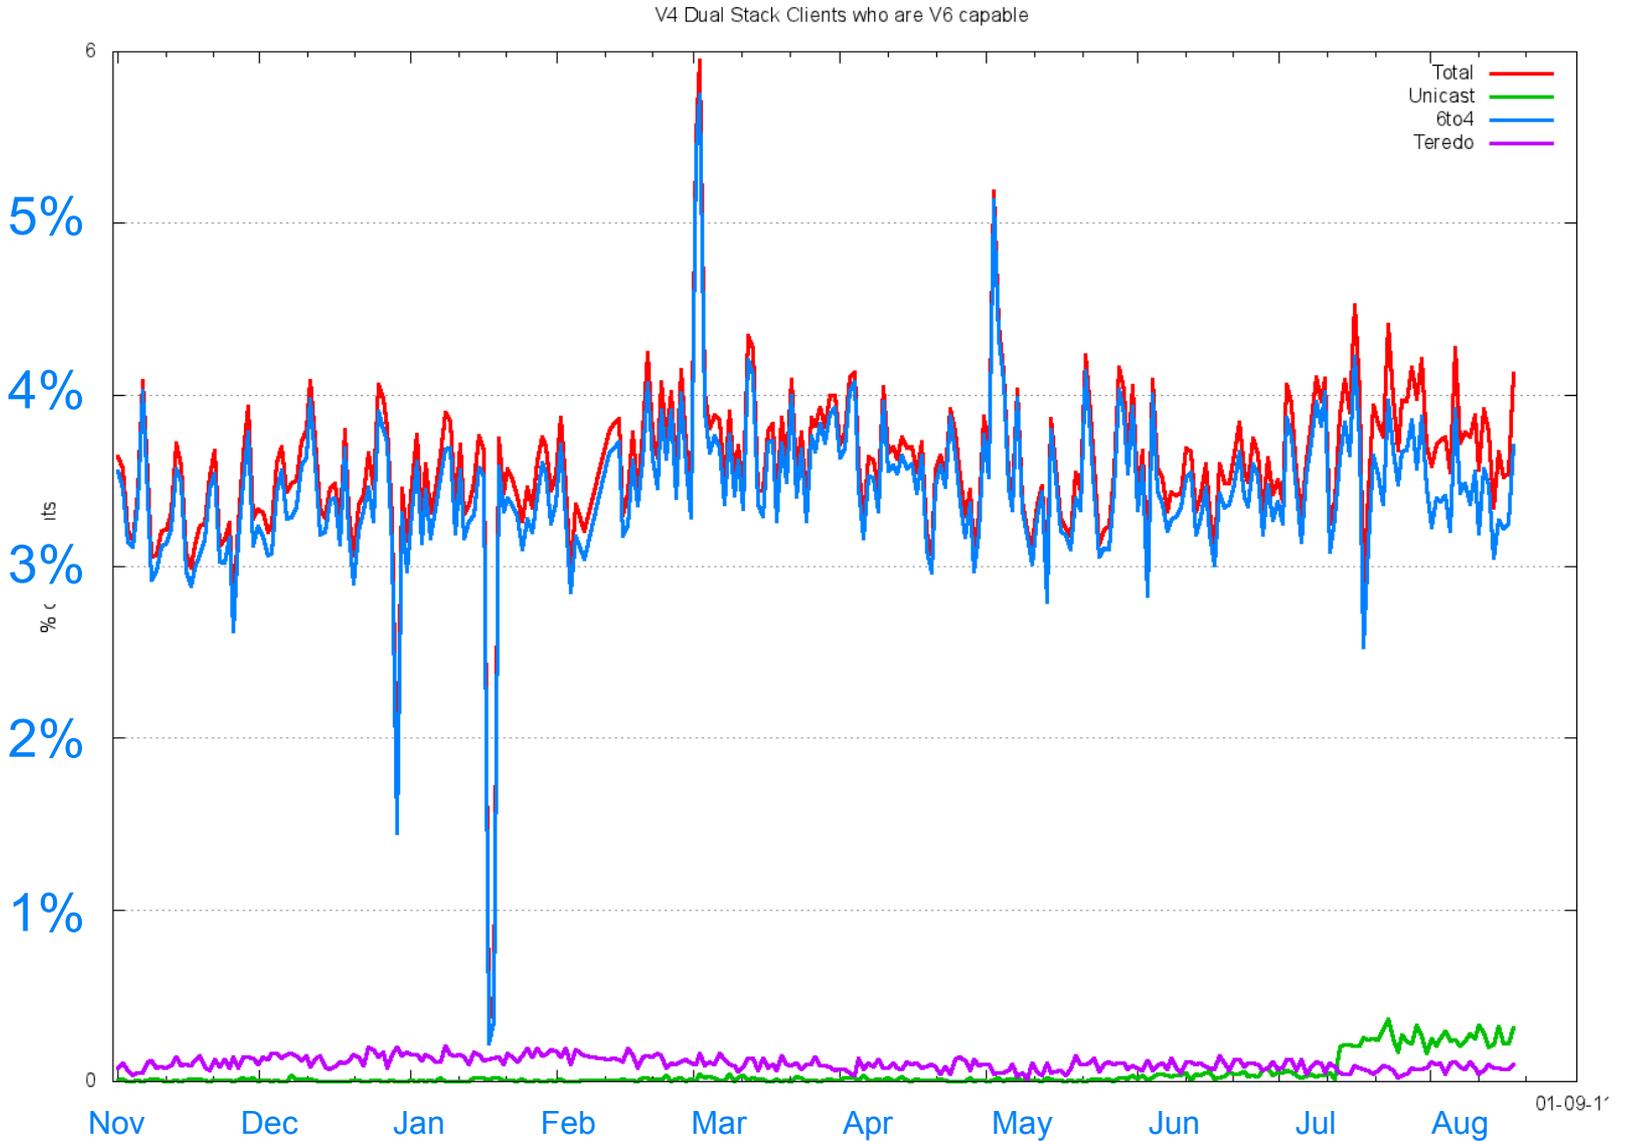
Client no Prefer V6 in Dual Stack by V6 Address Type



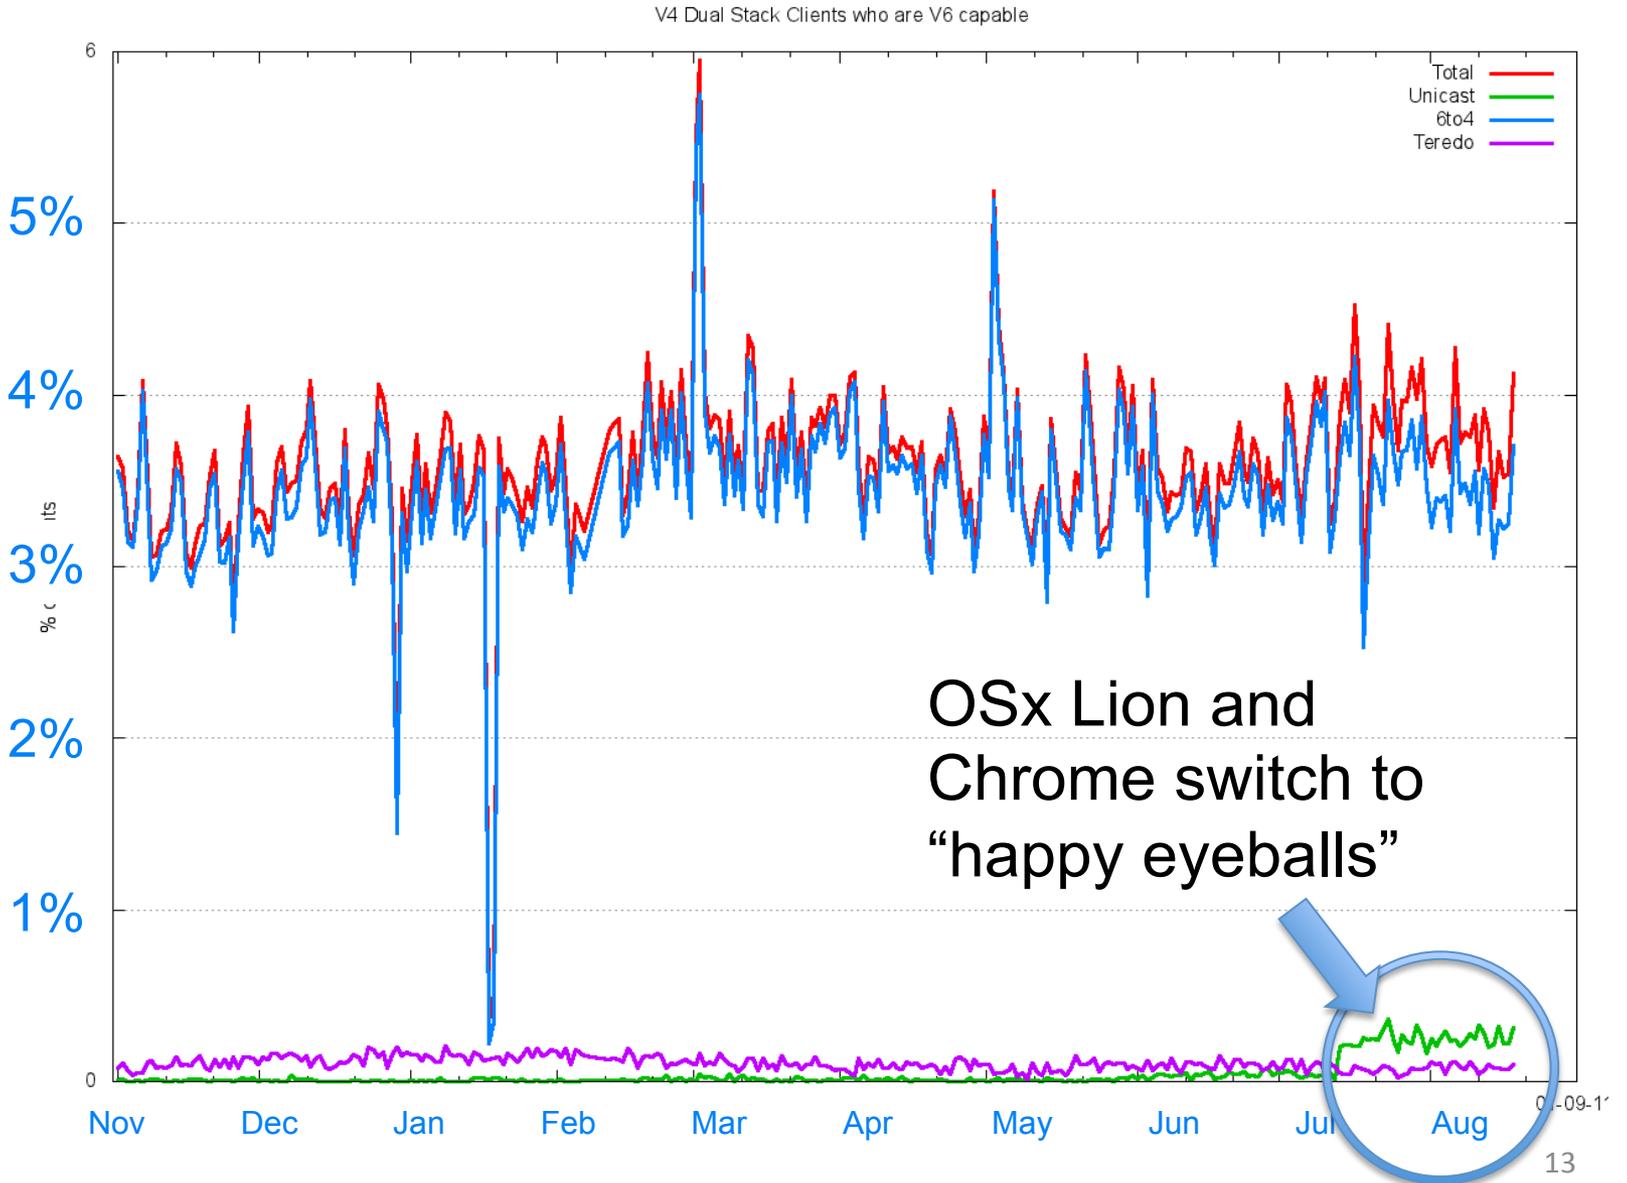
Is This All There Is?

- **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

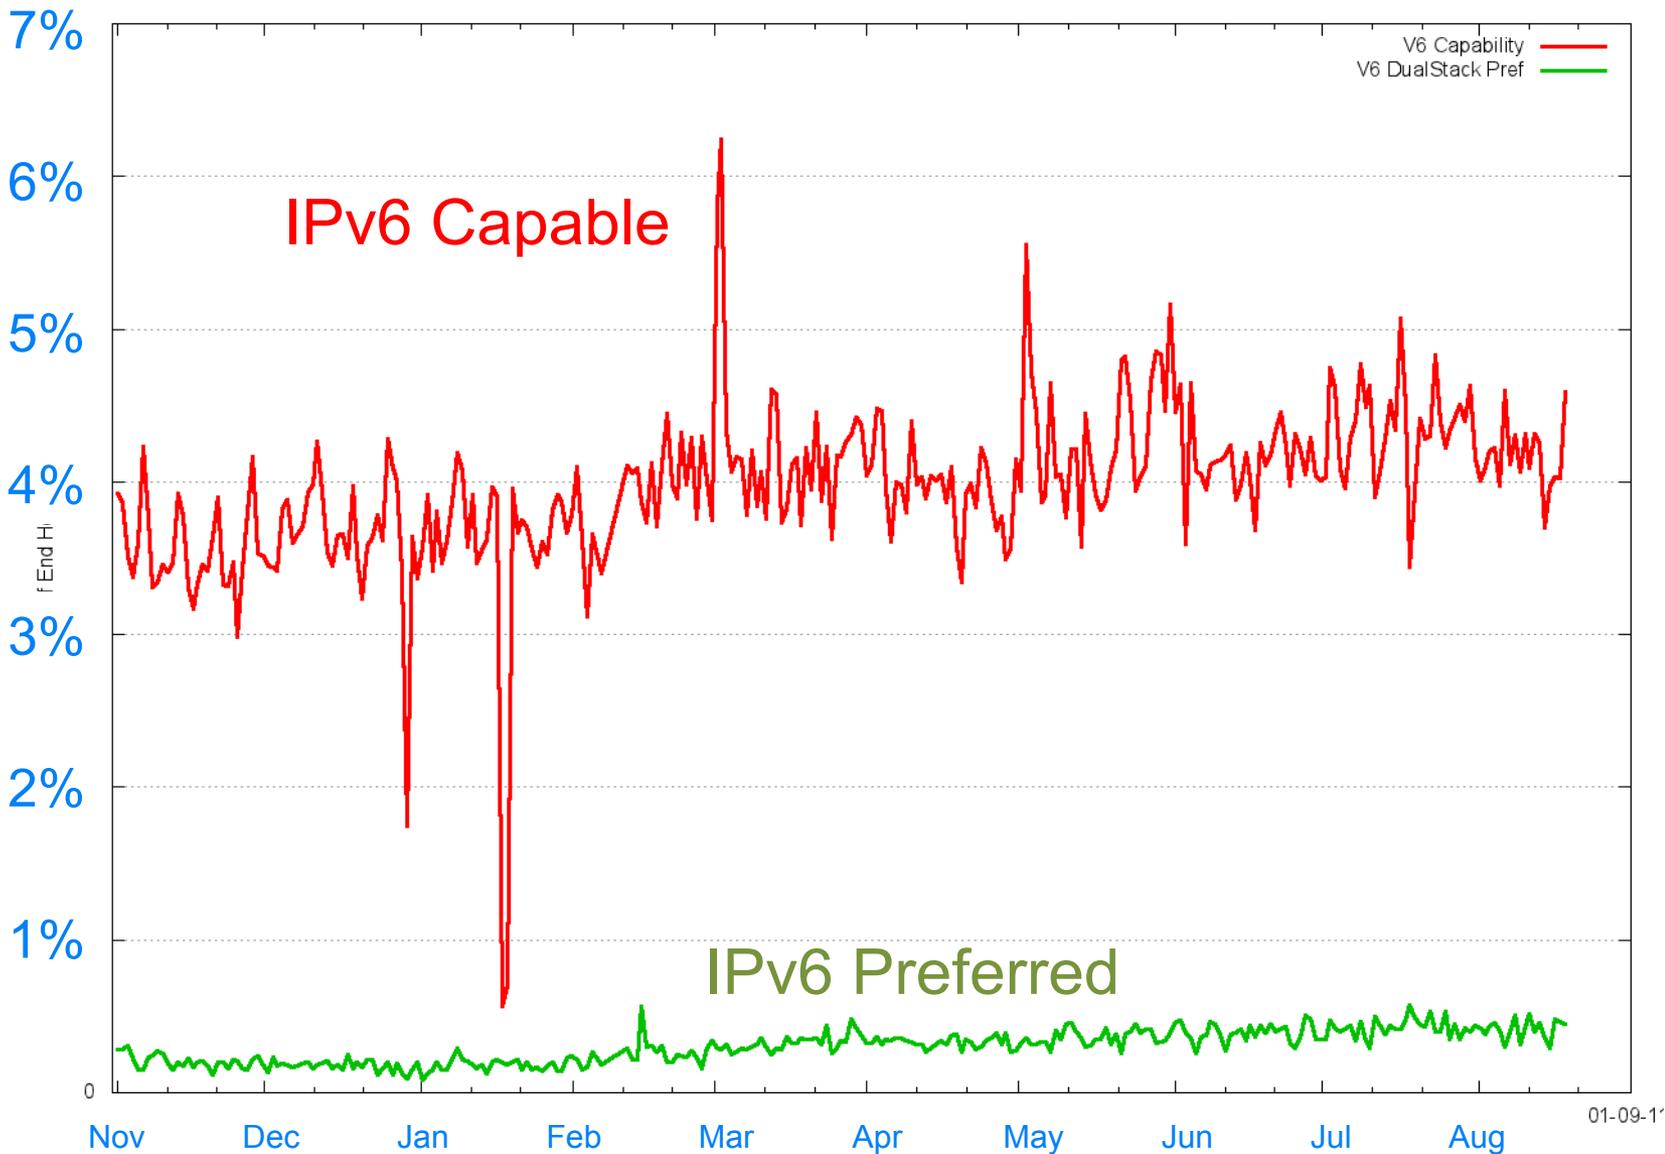


IPv6 Capable Clients



IPv6: "could" vs "will"

IPv6-Capable End Hosts



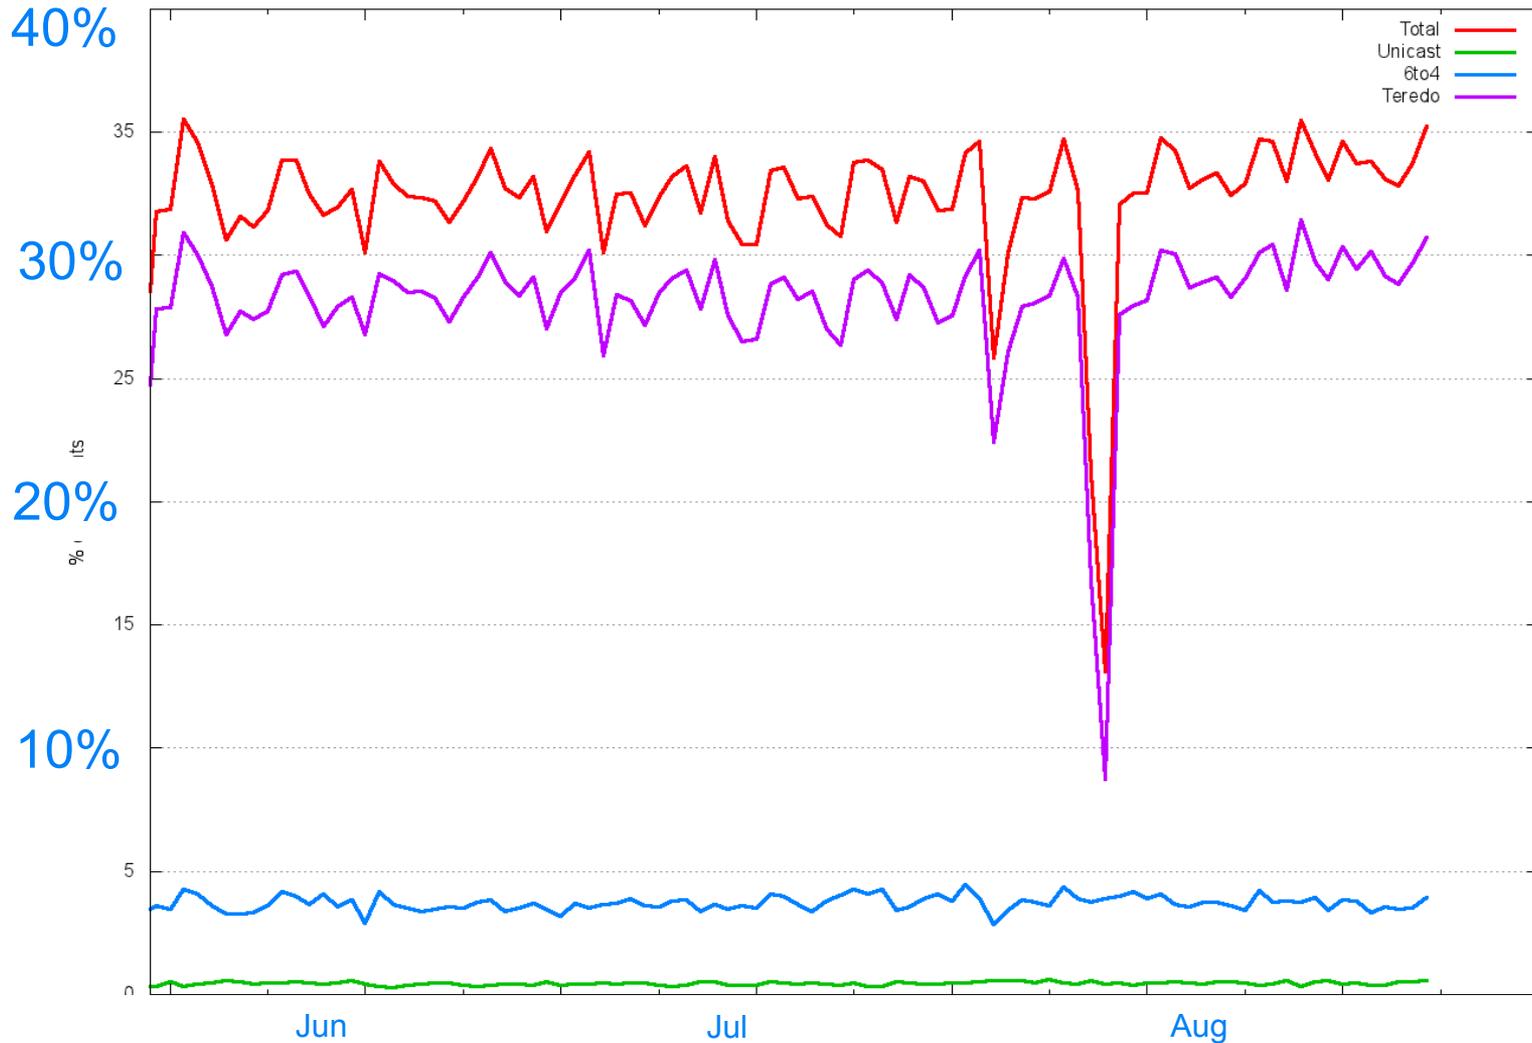
Is This All There Is?

- **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](http://[2401:2000:6660::f003]/1x1.png)

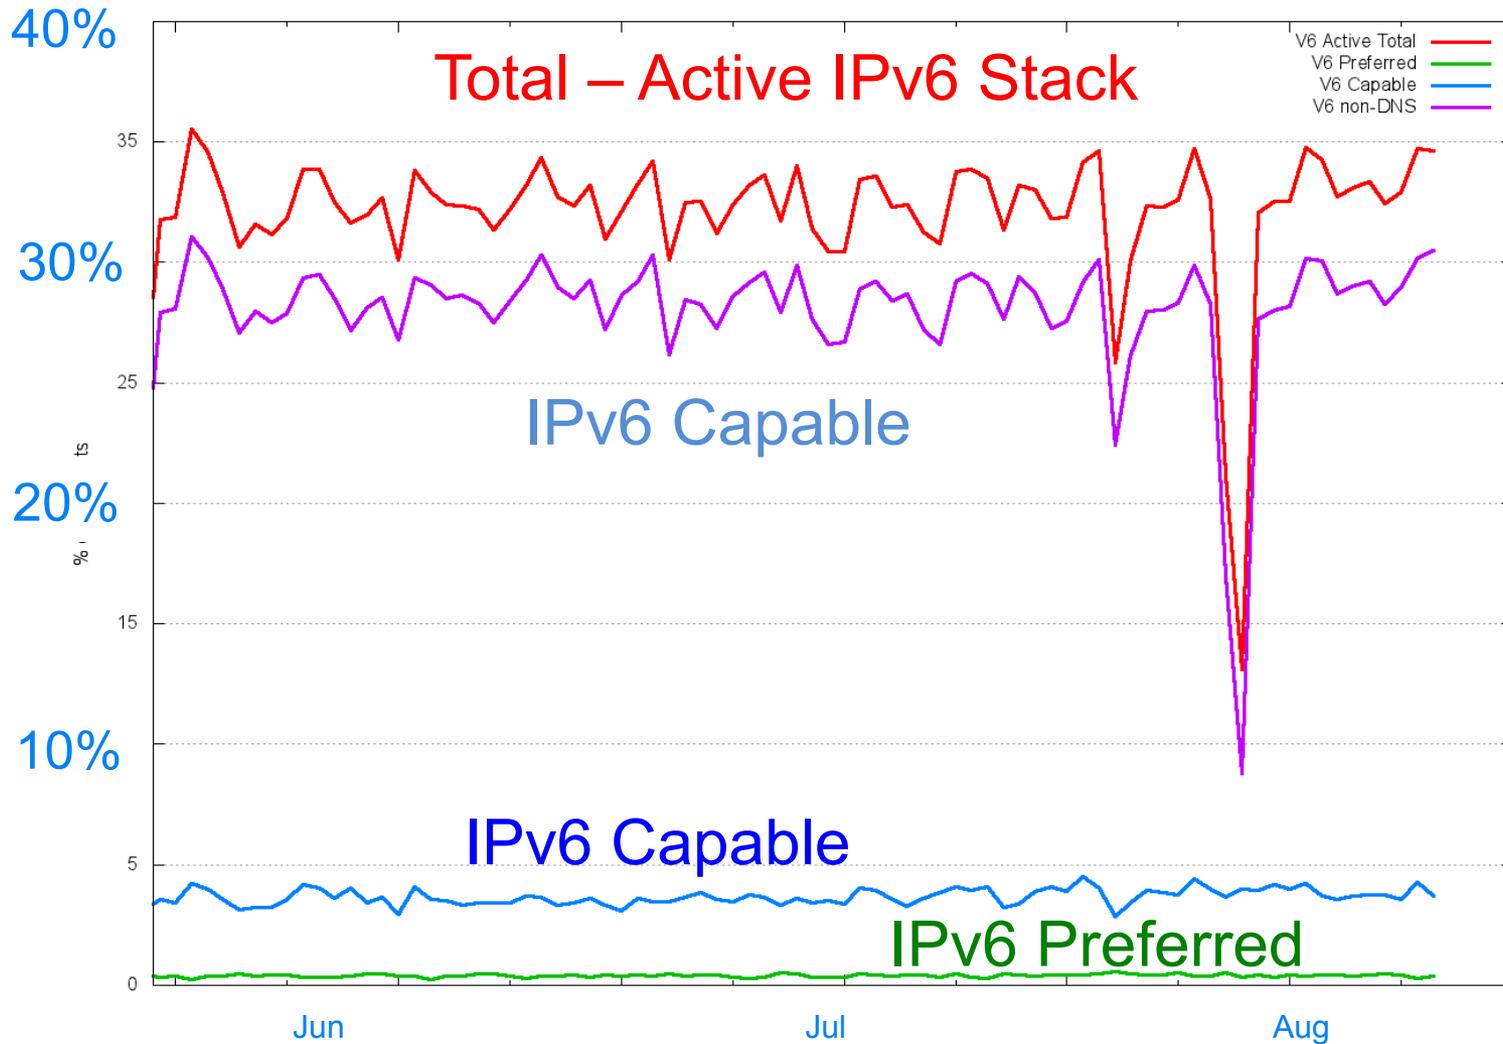
IPv6 "Coerceable" Clients

Clients who support V6 Literal by V6 Address Type



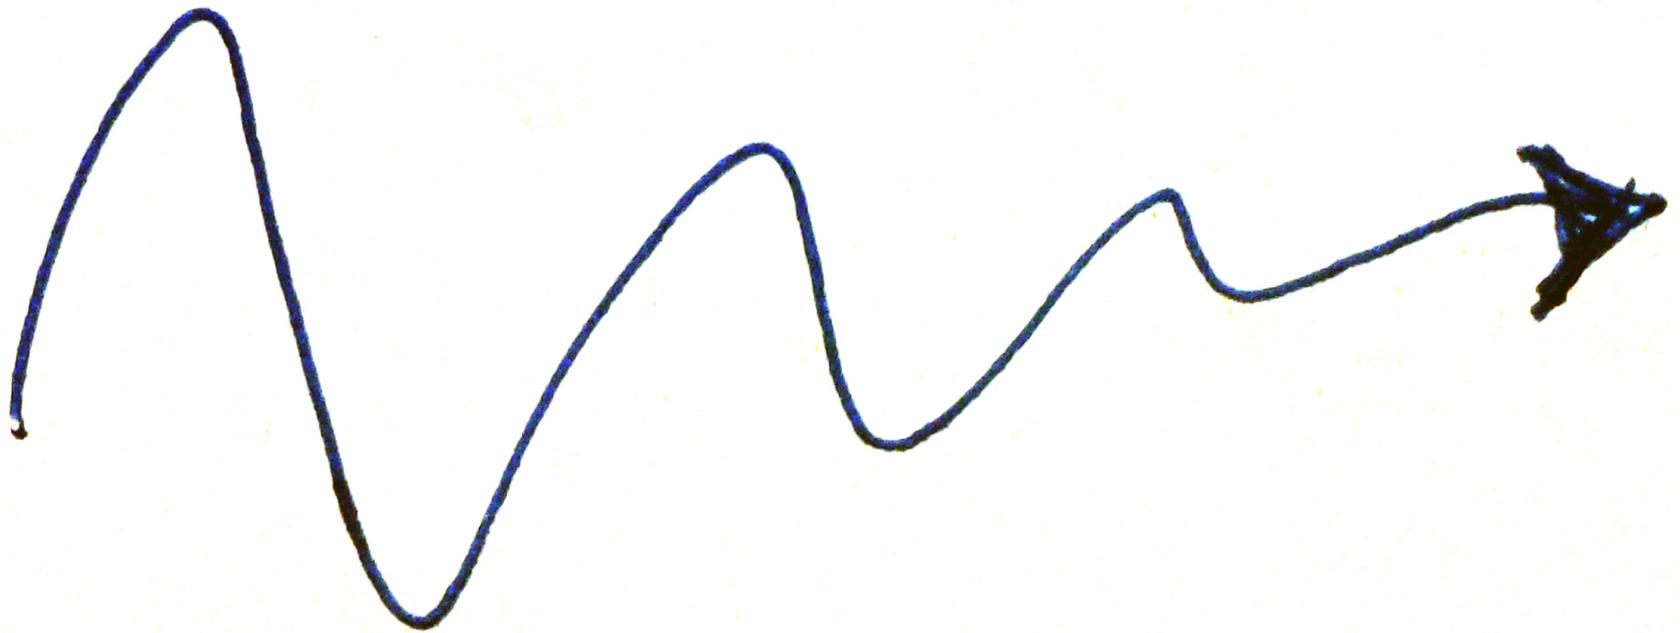
IPv6 Client Capabilities

Summary of Client V6 Capability



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



Some Measurements

39% of the IPv4 transit networks appear to be dual stack capable

48% of the Internet's end devices have an installed IPv6 stack that can be tickled into life

0.3% of the Internet's end devices have native IPv6 delivered to them

Some Measurements

39% of the IPv4 transit networks appear to be dual stack capable

48% of the Internet's end devices have an installed IPv6 stack that can be tickled into life

0.3% of the Internet's end devices have native IPv6 delivered to them

Where's the problem here?

The last mile access service business is not doing IPv6 because:

- A) they are stupid
- B) they are lazy
- C) they are uninformed
- D) they are broke
- E) they operate in an economic and business regime that makes provisioning IPv6 an unattractive investment option for them

The last mile access service business is not doing IPv6 because:

A) they are stupid

B) they are lazy

C) they are uninformed

D) they are broke

 **E)** they operate in an economic and business regime that makes provisioning IPv6 an unattractive investment option for them

Hint!

The last mile access service business is not doing IPv6 because:

A) they are stupid

B) they are lazy

C) they are uninformed

D) they are broke

Why is IPv6 such an unattractive business proposition for Carriage Providers?

They operate in an economic and business regime that makes

provisioning IPv6 an unattractive investment option for them

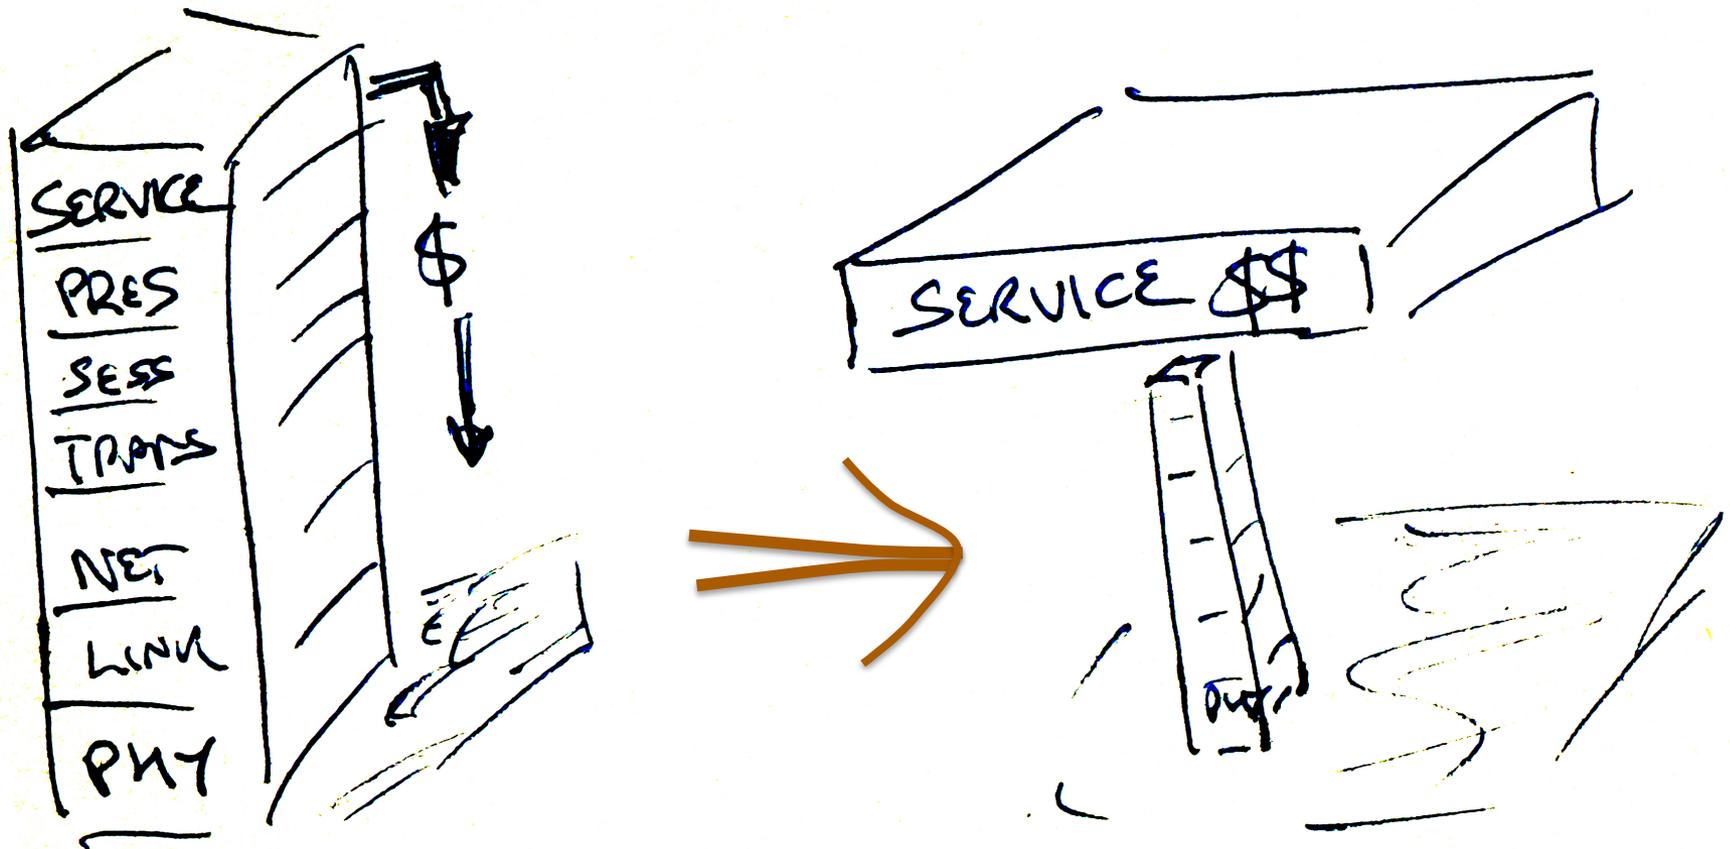
Back to networking basics....



Telco nostalgia...

The historical
vertically integrated
service architecture

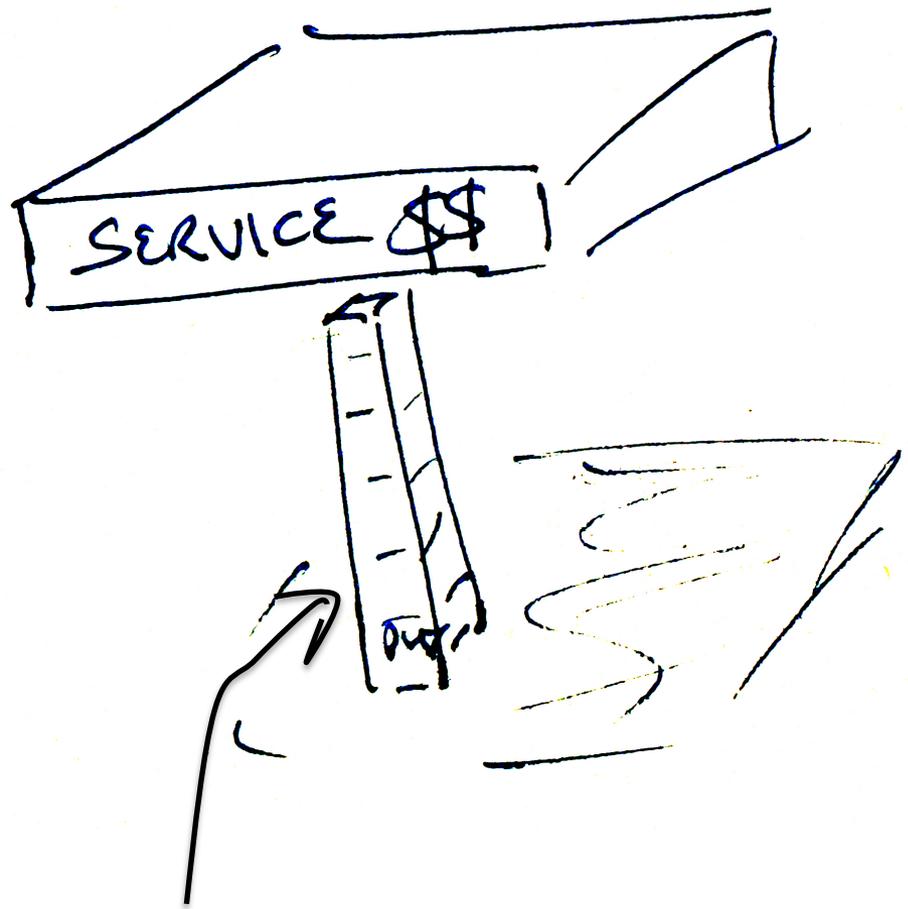




Devolution of the integrated service architecture through an open IP service architecture and deregulation



Devolution of the integrated
service architecture

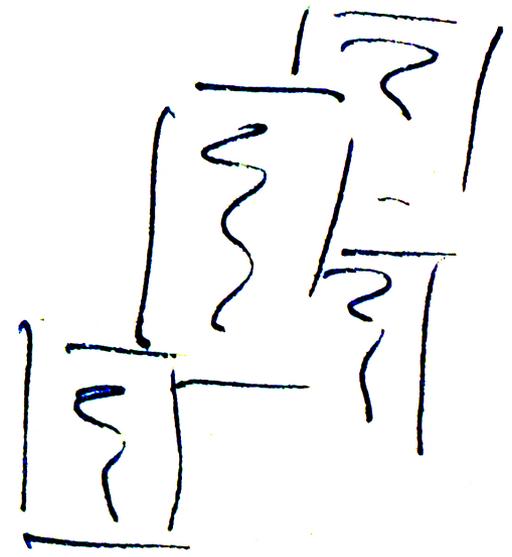


Where's the money to invest
in new network services?

Users



Services



Access Provider

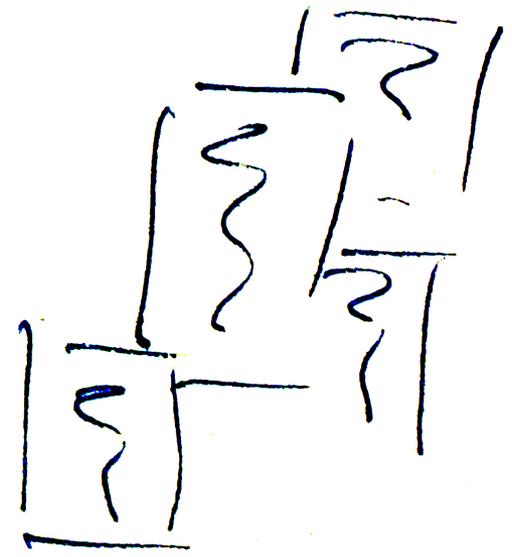


Users



Access Provider

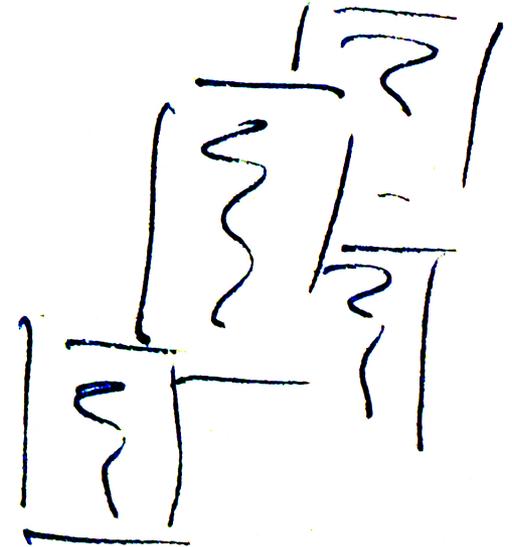
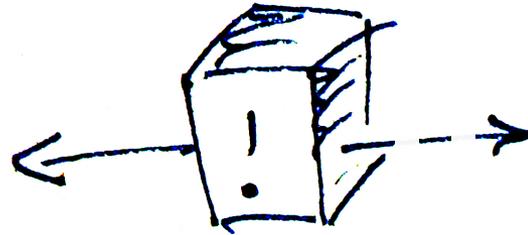
Services



CGNs and ALGs and similar IPv4
rationing middleware devices
provide control points in the IPv4
network that allow monetary
extraction from both consumers and
content providers

Users

Services



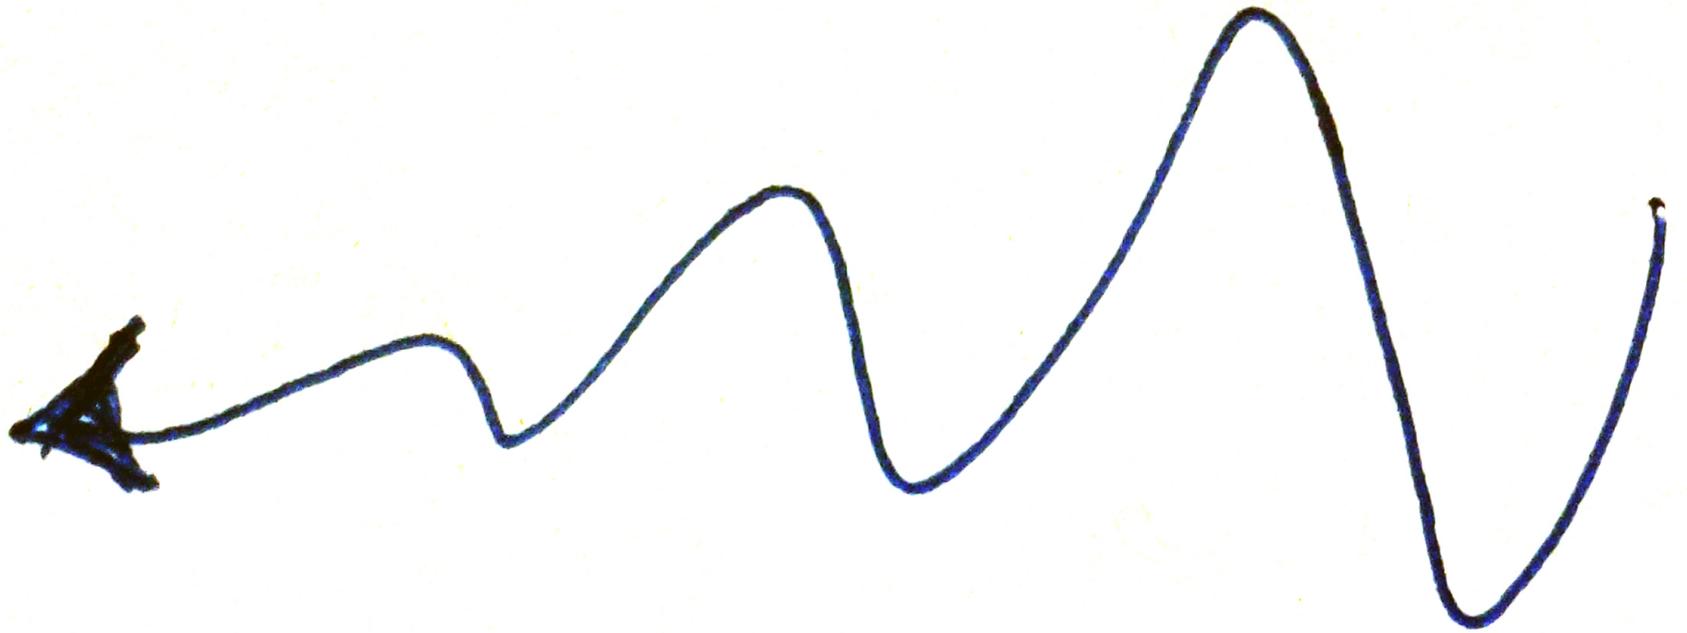
Access Provider



A digression...

How "real" is this
risk?

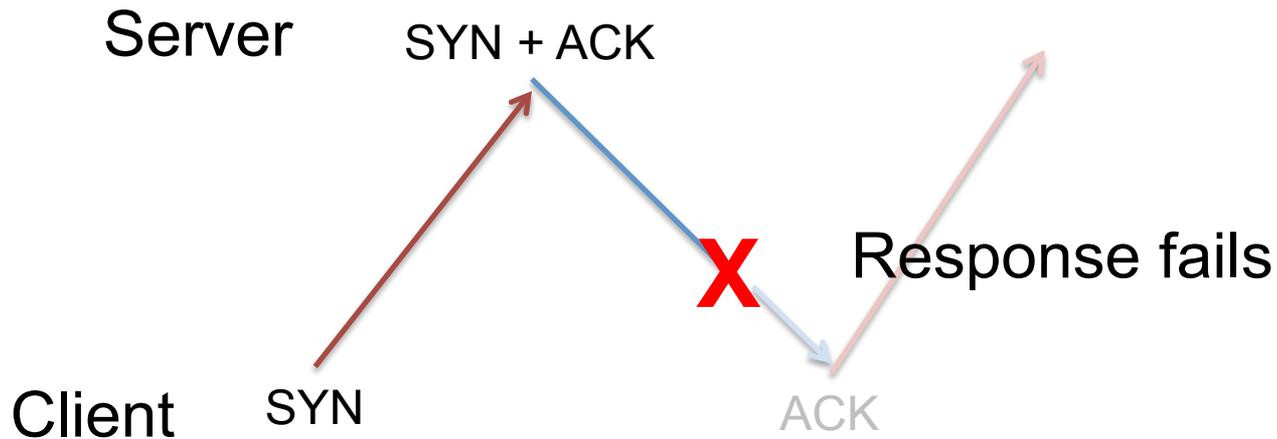
Is this industry seriously prepared to contemplate an
IPV4 forever strategy? **Yes — it's a possibility!**



Failure Observations

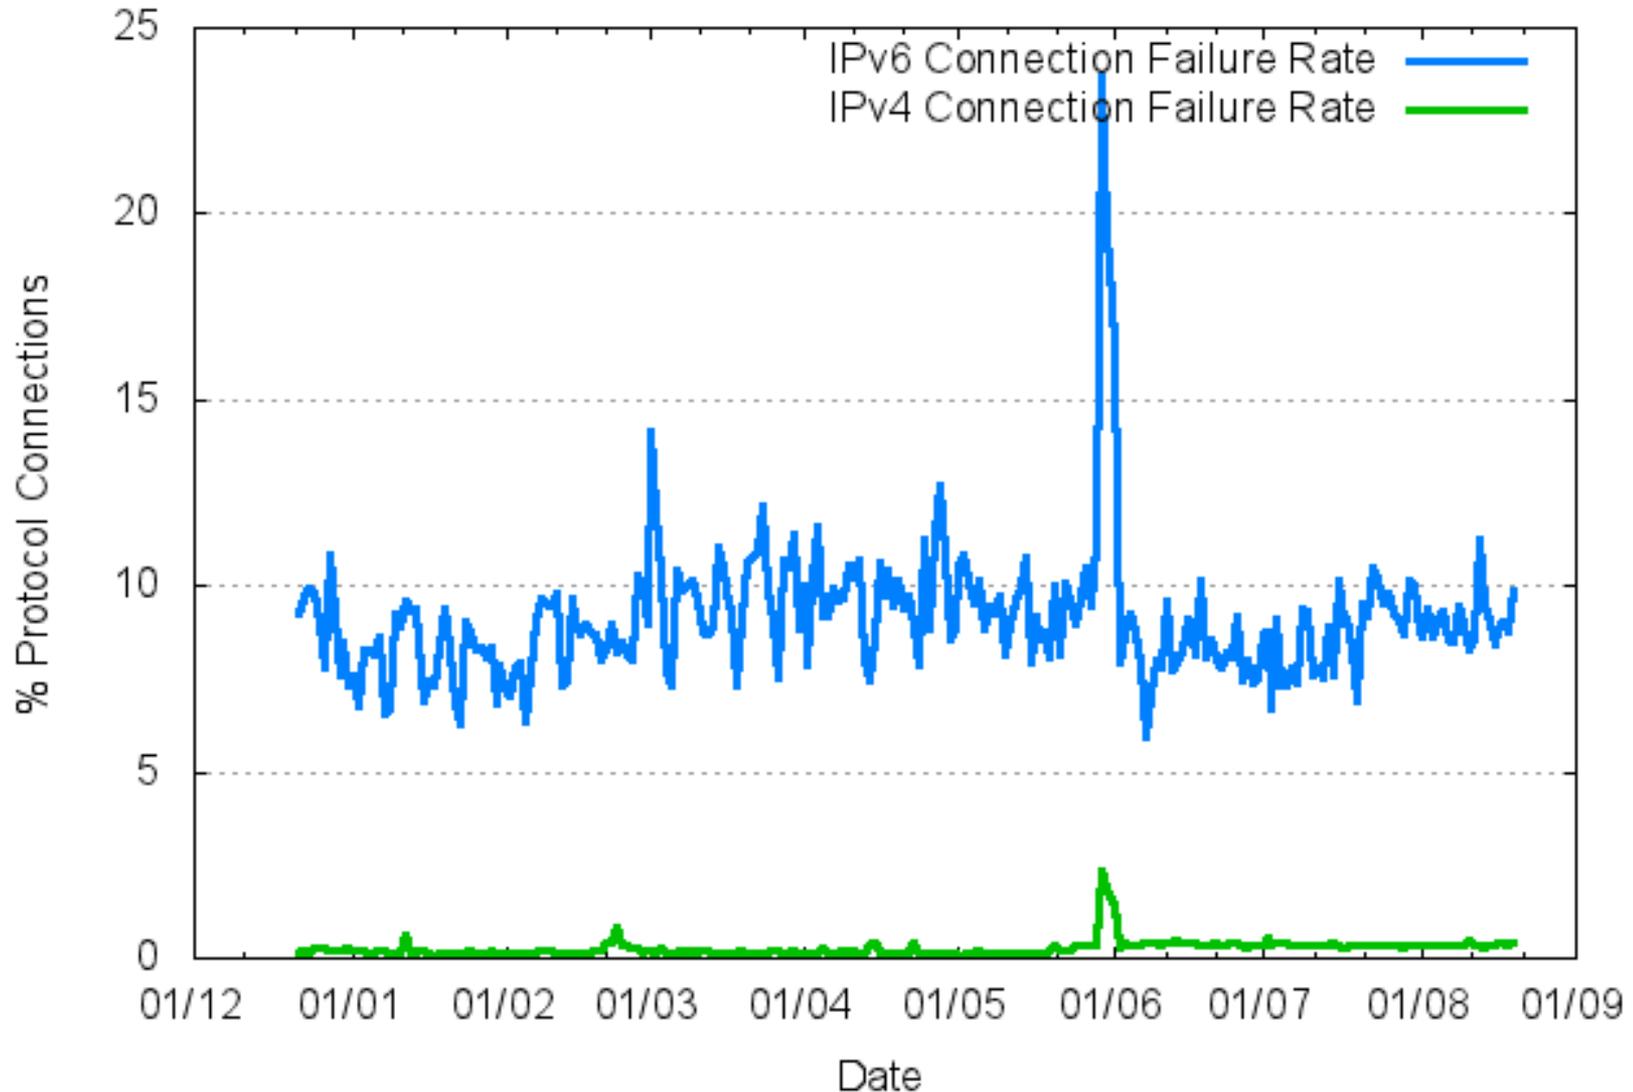
Connection Failure

To attempt to look at some instances of connection failure, let's look for connections that fail after the initial TCP SYN

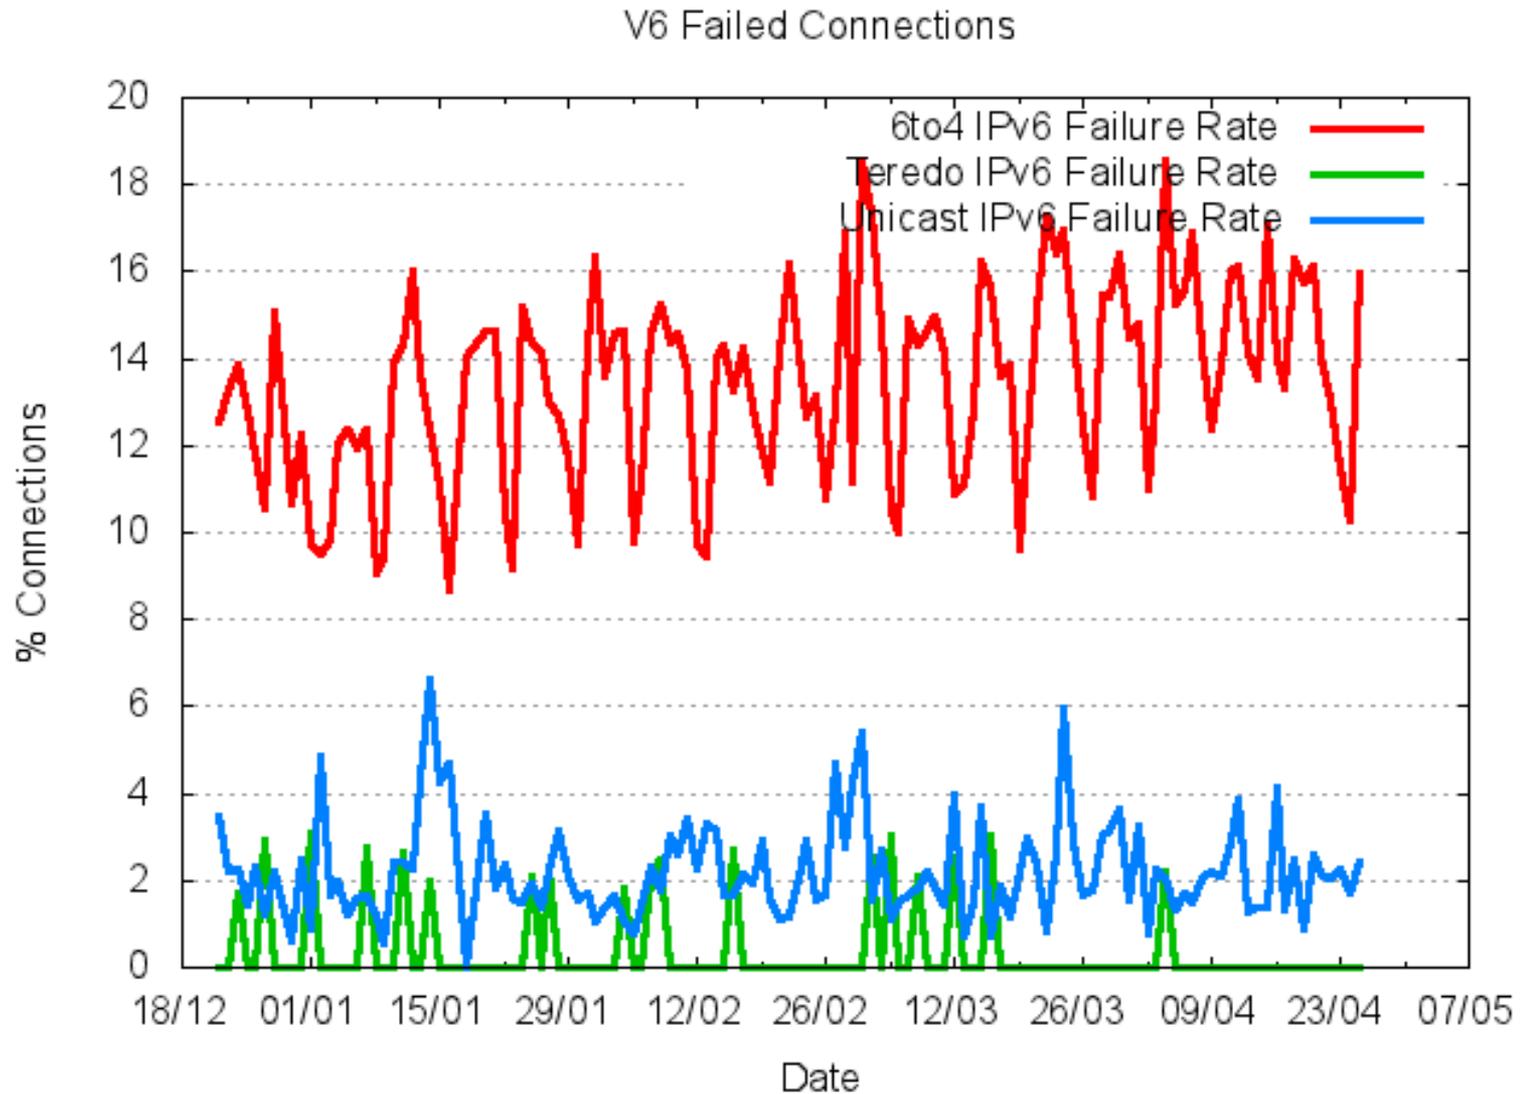


Connection Failure

Relative Percentage of Failed Connections



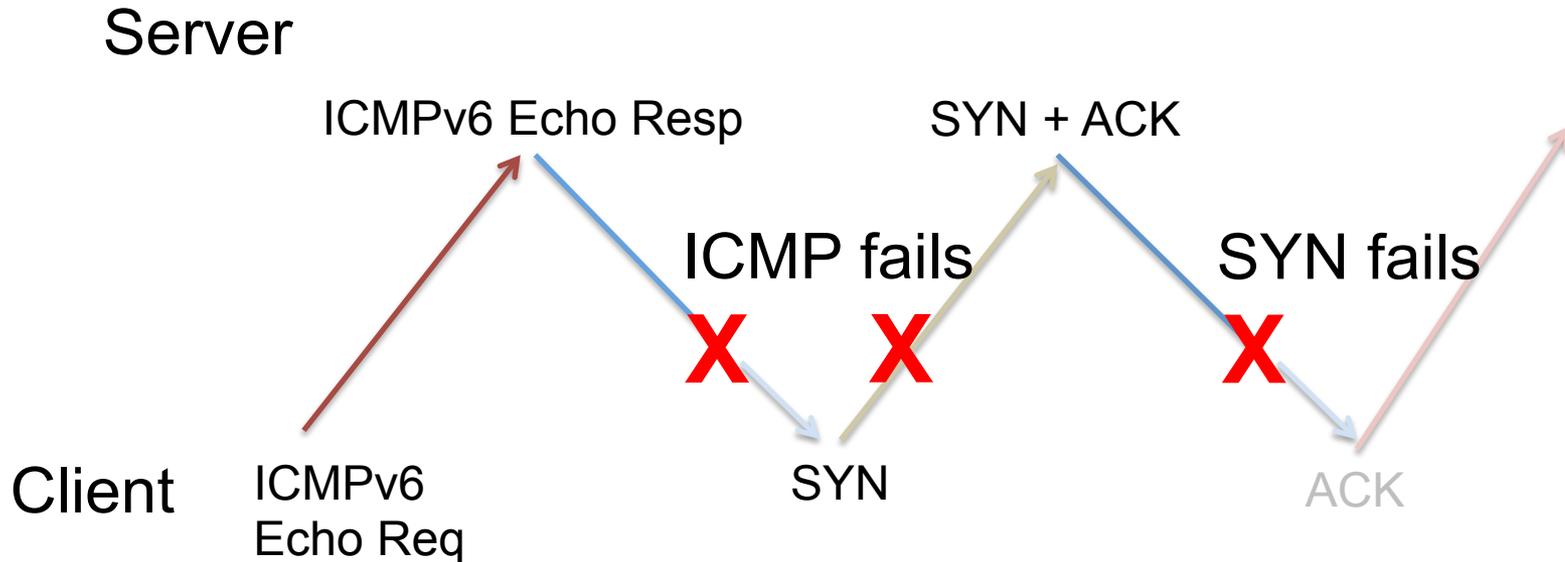
IPv6 Connection Failure



**Is Teredo really THAT
good?**

Teredo Connection Failure

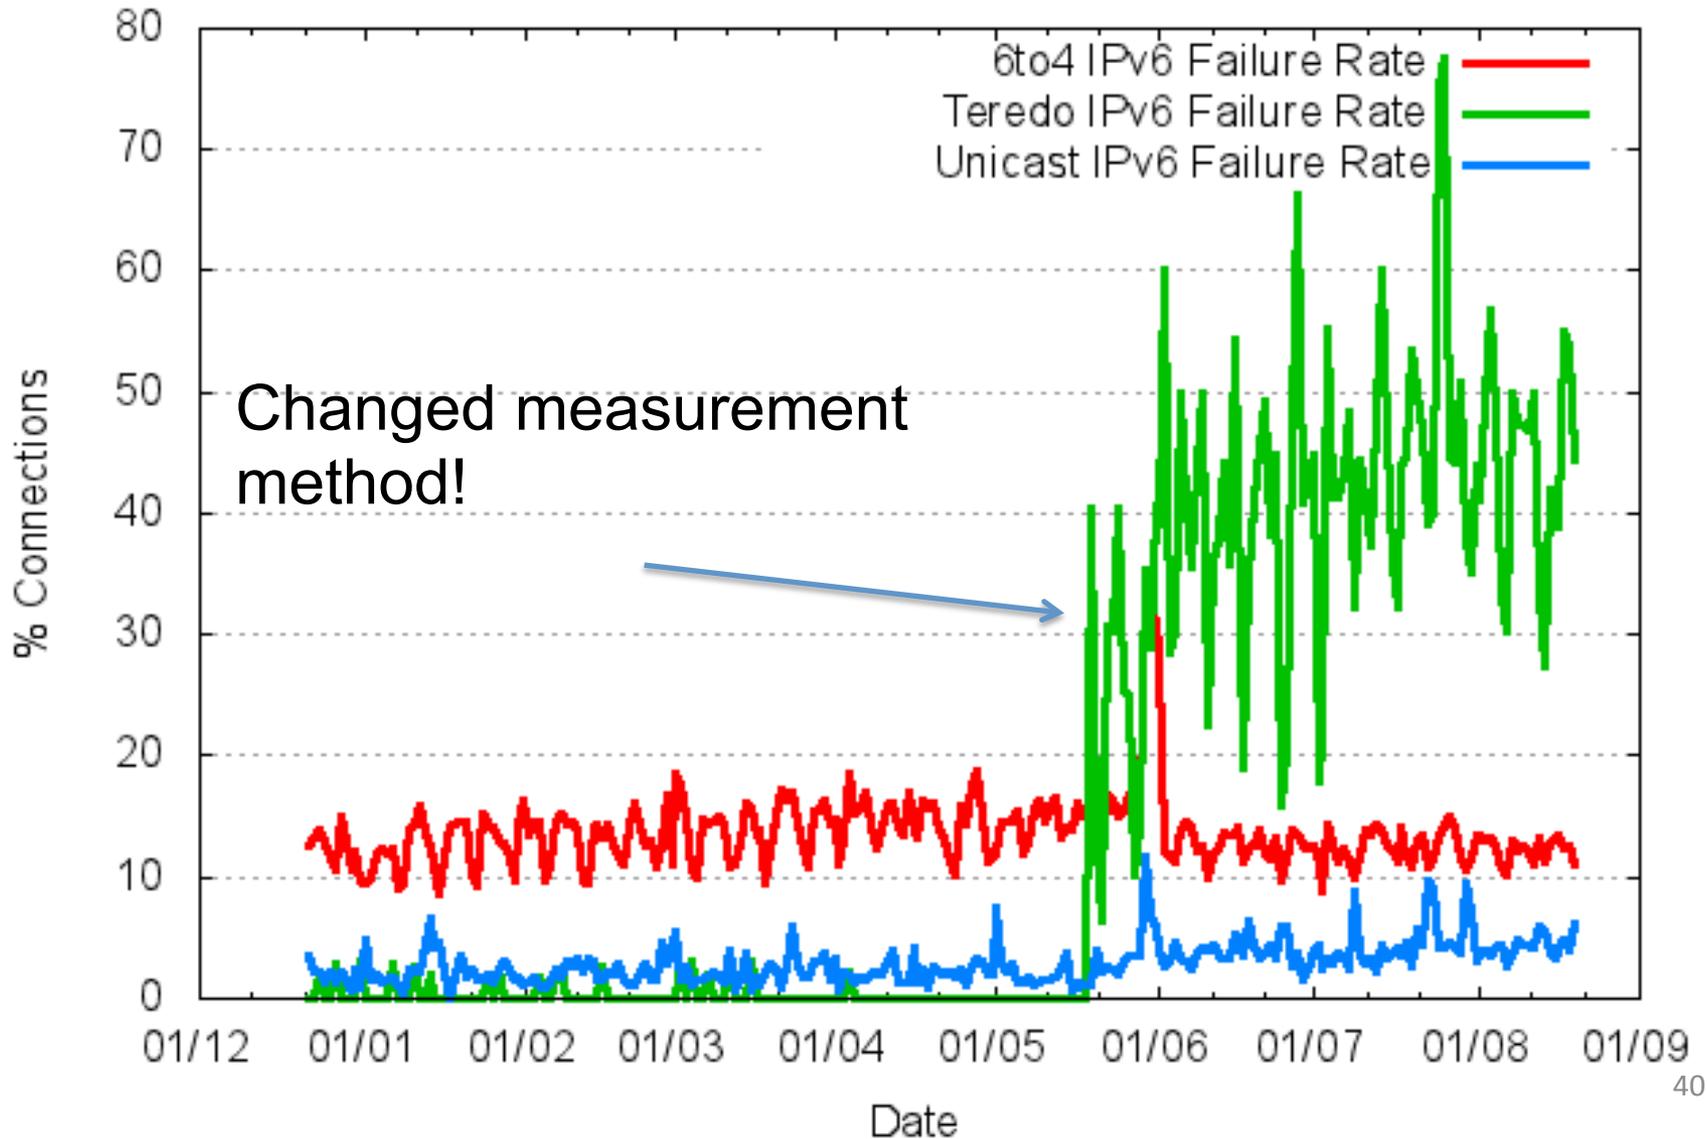
Teredo uses an initial ICMPv6 exchange to assist in the Teredo Server / Relay state setup



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

V6 Failed Connections



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!
 - And its not local firewall rules!
 - Teredo's NAT traversal is failing 45% of the time

Teredo's NAT traversal
algorithm is failing 45% of
the time

What have we learned about applications and their ability to perform NAT traversal for multi-party NAT bindings?

This is seriously broken!

NATs are incredibly difficult and
unreliable for applications to
cope with!

What about CGNs?

CGNs are just big remote NATs

What can we say about applications and CGN traversal for multi-party NAT bindings?



Thank You!

