



# Measuring IPv6

Geoff Huston

George Michaleson

APNIC Labs, May 2013



# What's the question?

## The Big Question:

How "well" are we going with the transition to IPv6?



# What's the question?

## The Big Question:

How "well" are we going with the transition to IPv6?

That's a very difficult question to measure!



# « Measurable » Questions

- How much traffic uses IPv6?
- How many connections use IPv6?
- How many routes are IPv6 routes?
- How many service providers offer IPv6?
- How many domain names have AAAA RRs?
- How many domain NS's use AAAA's?
- How many DNS queries are for AAAA RRs?
- How many DNS queries are made over IPv6?
- How many end devices have IPv6?
- How many end devices use IPv6?

'''



# Back to the Big Question

- None of these specific measurement questions really embrace the larger question
- They are all aimed at measuring IPv6 within particular facets of the network infrastructure, but they don't encompass all of the infrastructure of the network at once



# Back to the Big Question

- To make an IPv6 connection everything else (routing, forwarding, DNS, transport) has to work with IPv6
- So can we measure how many connected devices on today's Internet are capable of making IPv6 connections?



# An Observation...

The conventional view of transition was that end hosts would use a very simple protocol selection algorithm:

- If the local host has an IPv6 interface, and the remote host has an IPv6 address, then always try to connect using IPv6.
- Otherwise use IPv4.



# How to measure IPv6 in the Internet

- Set up a service on both IPv6 and IPv4
- Measure the proportion of users who connect to the service using IPv6



# But...

- We tried this on <http://www.apnic.net> in 2010
  - We found a very high number of IPv6 users (~5%)
  - Why?
  - Small, geek-centric client population of users of this service have biased the measurement!
- We really need to use a massively popular web service to conduct this experiment
  - But “massively popular web services” worry constantly about service resiliency and privacy of their data regarding users
  - So they tend to be extremely suspicious of adding Javascript elements to their service that performs third party dual stack tests with their clients (and I can't blame them!)
- So we need to rethink this approach...



How to measure a million end users for  
their IPv6 capability



# How to measure a million end users for their IPv6 capability

- Be Google (or any other massively popular web service provider)



# How to measure a million end users for their IPv6 capability

- Be Google (or any other massively popular web service provider)

or



# How to measure a million end users for their IPv6 capability

- Be Google (or any other massively popular web service provider)

or

- Get your code to run on a million users' machines



# Ads are ubiquitous

**REMINDER:**  
SOMETIMES YOU  
NEED TO LET THE  
WILD OUT  
(remember to breathe)



should not profit from region's name

80 comments

## Cutting cord too early 'risks health'



**Exclusive:** Childbirth experts query policy after research suggests early clamping of umbilical cord can lead to iron deficiency anaemia

46 comments

● Mother sings praises of delayed clamping

## Chinese official sacked for excess



Communist boss in Jiangsu province begs in vain for forgiveness after campaigners gatecrash lavish dinner

17 comments

## Measles cases rise to 942 in Wales



Figure for greater Swansea area rises by 56 as experts warn epidemic shows no sign of easing

- Big drive to halt measles outbreak
- Measles vaccination campaign begins
- Outbreak triggers fresh emphasis on vaccination
- The story behind the MMR scare
- Measles and MMR: the essential guide

## PM handed press regulation dilemma



Cross-party plans rejected as papers launch audacious bid to set up own royal charter-backed body

197 comments

- Read the draft alternative royal charter
- Alternative regulation plans: the key differences
- Editorial: time for a ceasefire


## Ukip election candidate suspended



Antisemitic comments were allegedly posted on conspiracy theory website under Anna-Marie Crampton's name but she says she is hacking victim

- Farage: Ukip candidates may have BNP past
- Clegg kills 'snooper's charter' bill
- Nick Thornsby: Clegg reminded he is a liberal

## 10 of the worst



**George Monbiot**  
My search for a smartphone that isn't soaked in blood



**Spare Rib**  
Back for more



**Box set gold**  
Big Train



**Measles & MMR**  
Essential guide



**Turner prize**

## Ballads of a thin man

★★★★★



Iggy and the Stooges can still make a racket, but the best songs on Ready to Die are the ballads, writes Alexis Petridis

17 comments

on a  
**Low Rate Credit Card**



with an ongoing purchase rate of 13.49% p.a. (variable).

Apply now

More Extra offers

Today's paper

The Guardian

G2 features

Comment and debate

Editorials, letters and corrections

Obituaries

Other lives

Sport

Film & music

Subscribe

Vote for the Guardian



Contact us

How to contact the Guardian and Observer

Guardian readers' editor

Observer readers' editor

On this site

A-Z

Blogs

Cartoons

Community

Corrections

Crosswords

Digital archive

Digital edition

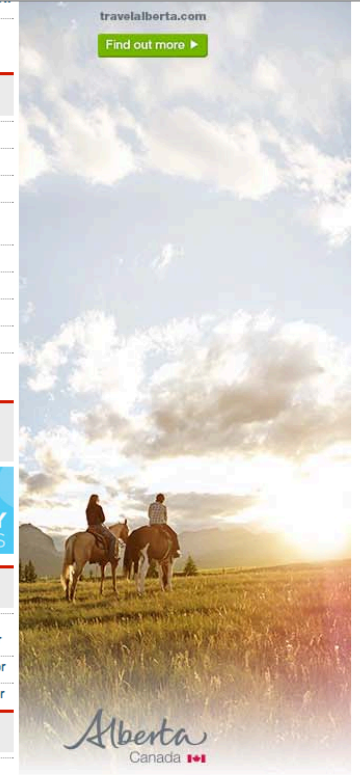
G24

guardian.co.uk in 1821

Guardian mobile

travelalberta.com

Find out more



Alberta  
Canada

## The price of resistance in DRC

Plagued by an armed militia, villagers in the Democratic Republic of the Congo have fought back - but at a cost



AC Jimbo's European papers review



# Ads are ubiquitous

The image shows a screenshot of a news website with a red circle highlighting several areas. On the left, a large banner features a landscape with people on horseback and the text: "REMINDER: SOMETIMES YOU NEED TO LET THE WILD OUT (remember to breathe)". The main content area includes several news articles with images and headlines, such as "Cutting cord too early 'risks health'", "Chinese official sacked for excess", "Measles cases rise to 942 in Wales", "M handed press regulation dilemma", and "Ukip election candidate suspended". A prominent advertisement for a "Low Rate Credit Card" is highlighted in the center, showing a credit card and the text "Apply now". Other ads include "Ballads of a thimble" and "George Monbiot My search for a smartphone that isn't soaked in blood". The right sidebar contains a navigation menu with categories like "More Extra offers", "Today's paper", "Comments and debate", "Editorials, letters and corrections", "Opinion", "Other news", "Sport", "Film & music", "Subscriptions", and "Vote for the Guardian". At the bottom right, there is a "travelalberta.com" banner with a "Find out more" button and the "Alberta Canada" logo.



Ads are ubiquitous



**DigiCert SSL Certificates** ✕

The Strongest Encryption Available, in a Variety of Options. Buy Now!

[www.DigiCert.com](http://www.DigiCert.com)

Ads by Google

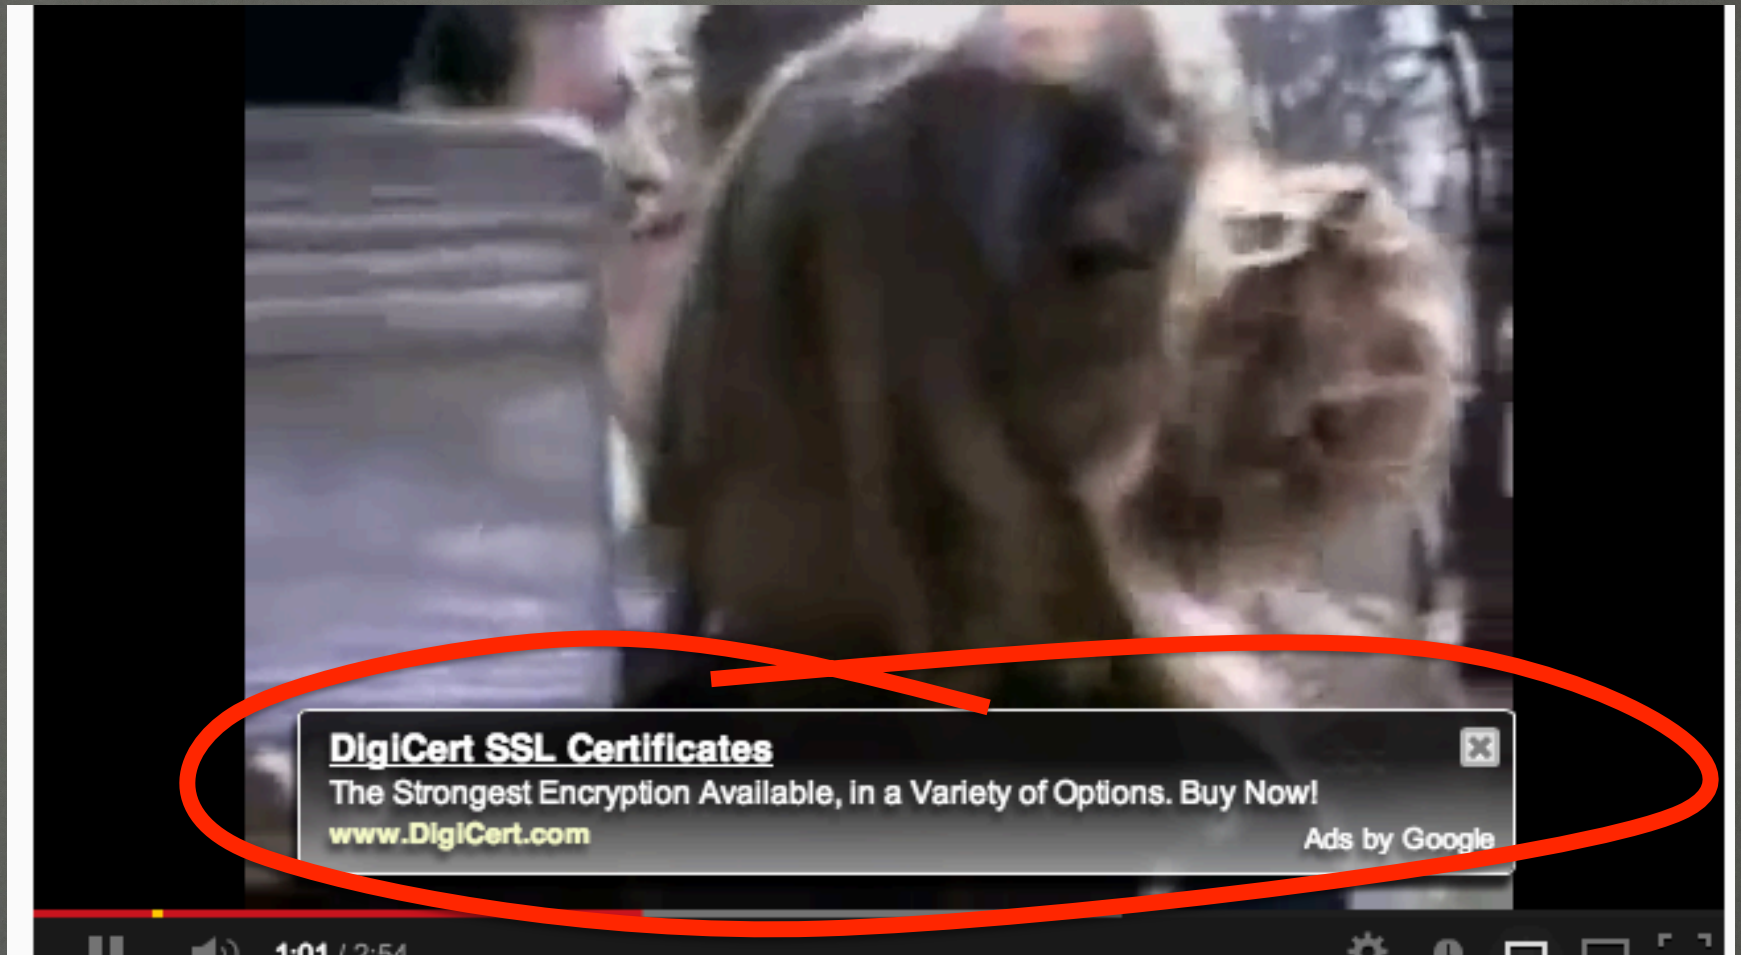


1:01 / 2:54





Ads are ubiquitous





# Ads are implemented in Adobe Flash

- Advertising channels use Flash to make ads interactive
  - This is not just an 'animated gif'

**0%**  
p.a.  
on purchases

\*New cards only.  
Conditions apply.

Apply now



on a  
**Low Rate Credit Card**



with an ongoing  
purchase rate of  
13.49% p.a. (variable).

Apply now





# Flash makes ads interactive

- [Apply Now] hover-over is interactive, and responds when selected.

0%  
p.a.  
on purchases

\*New cards only.  
Conditions apply.

Apply now



on a  
Low Rate Credit Card



with an ongoing  
purchase rate of  
13.49% p.a. (variable).

Apply now





# Flash and the network

- Flash includes primitives in 'actionscript' to fetch 'network assets'
  - Typically used to load alternate images, sequences
  - Not a generalized network stack, subject to constraints:
    - Port 80
    - crossdomain.xml on hosting site must match source name (wildcard syntax)
- Flash has asynchronous 'threads' model for event driven, sprite animation



# APNIC's measurement technique

- Craft flash/actionscript which fetches network assets to measure.
- Assets are reduced to a notional 'x' image which is not added to the DOM and is not displayed
- Assets can be named (gethostbyname()) or use literals (bypass DNS based constraints)
- Encode data in the name of fetched assets
  - Result is returned by DNS name with wildcard



# Advertising placement logic

- Fresh Eyeballs == Unique IPs
  - We have good evidence the advertising channel is able to sustain a constant supply of unique IP addresses
- Pay by clicks, or pay by impression
  - If you select a preference for impressions, then the channel tries hard to present your ad to as many unique IPs as possible
- Time/Location/Context tuned
  - Can select for time of day, physical location or keyword contexts (for search-related ads)
  - But if you don't select, then placement is generalized
- Aim to fill budget
  - If you request \$100 of placement, a day, then inside 24h algorithm tries hard to even placement but in the end, will soak place your ad to achieve enough views, to bill you \$100



# Advertising placement logic

- Budget: \$100 per day, at \$1.00 'CPM' max
  - Clicks per millepressions: aim to pay no more than \$1 per click but pay up to \$1 for a thousand impressions
- Even distribution of ads throughout the day
- No constraint on location, time
- Outcome: 350,000 placements per day, on a mostly even placement model with end of day 'soak' to achieve budget goal



# Measuring IPv6 via Ads

- Use Flash code that is executed on ad impression
    - Client retrieves set of “tests” that use unique DNS labels from an ad-controller

(<http://drongo.randapnic.net/measureipv6id.cgi?advertID=99999>)
    - Client is given 5 URLs to load:
      - Dual Stack object
      - V4-only object
      - V6-only object
      - V6 literal address (no DNS needed)
      - Result reporting URL (10 second timer)
- All DNS is dual stack



# Why These Tests?

- Dual Stack URL
  - Which protocol will the client PREFER to use?
- V4 only URL
  - Control comparison (Reliability, RTT)
- V6 only URL
  - Is the client CAPABLE of using IPv6?
- V6 Literal URL
  - Does the client have an IPv6 stack at all?
- Result URL
  - Did the client keep the experiment running, or was it terminated early?



# Experiment Server config

- There are three servers, identically configured (US, Europe, Australia)
- Server runs Bind, Apache and tcpdump
- Experiment directs the client to the "closest" server (to reduce rtt-related timeouts) based on simple /8 map of client address to region



# Collected Data

- Per Server, Per Day:
  - http-access log  
(successfully completed fetches)
  - dns.log  
(incoming DNS queries)
  - Packet capture  
All packets



# Collected Data

## Web Logs:

```
h.labs.apnic.net 2002:524d:xxxx::524d:xxxx [29/Apr/2013:05:55:05 +0000] "GET /1x1.png?  
t10000.u7910203317.s1367214905.i888.v1794.v6lit  
h.labs.apnic.net 2002:524d:xxxx::524d:xxxx [29/Apr/2013:05:55:05 +0000] "GET /1x1.png?  
t10000.u7910203317.s1367214905.i888.v1794.r6.td  
h.labs.apnic.net 82.77.xxx.xxx [29/Apr/2013:05:55:05 +0000] "GET /1x1.png?  
t10000.u7910203317.s1367214905.i888.v1794.rd.td  
h.labs.apnic.net 82.77.xxx.xxx [29/Apr/2013:05:55:05 +0000] "GET /1x1.png?  
t10000.u7910203317.s1367214905.i888.v1794.r4.td  
h.labs.apnic.net 82.77.xxx.xxx [29/Apr/2013:05:55:05 +0000] "GET /1x1.png?  
t10000.u7910203317.s1367214905.i888.v1794&r=zrdtd-348.zr4td-376.zr6td-316.zv6lit-228
```

(In this case the client is using 6to4 to access IPv6, and prefers to use IPv4 in a dual stack context)



# Data Processing

- Web Logs:
  - V6 Capable/Preferred host counts
  - Breakdown of Teredo/6to4 vs Unicast
- Packet Logs:
  - Connection Failure counts (incomplete TCP handshake)
  - Performance measurements (TCP RTT)



# Reports

We perform a basic scan of the daily data and produce a number of reports:

a) A “summary” report of capabilities

<http://www.potaroo.net/ipv6/>

## IPv6 Measurements


This page contains an index of reports generated by the IPv6 capability test script. A description of the script itself and the statistics gathered by the script in April 2010 can be found in the article "Measuring More IPv6" (<http://www.potaroo.net/ispcol/2010-04/ipv6-measure.html>).

An associated [report](#) contains data that examines the connection failure rate recorded as part of this measurement activity.

### V6 Capability and Dual Stack Preference



V6 End hosts - comparison of V6 capability vs Dual Stack preference for V6



V6 End hosts - comparison of V6 only hosts vs Dual Stack preference for V6



V6 End hosts - comparison of hosts with Dual Stack V6 DNS capability vs Dual Stack V6 TCP capability vs V6 only



V6 Capable End Hosts by V6 Address Type



Dual Stack V6 Preferred V6 Capable End Hosts by V6 Address Type



Dual Stack V4 Preferred V6 Capable End Hosts by V6 Address Type

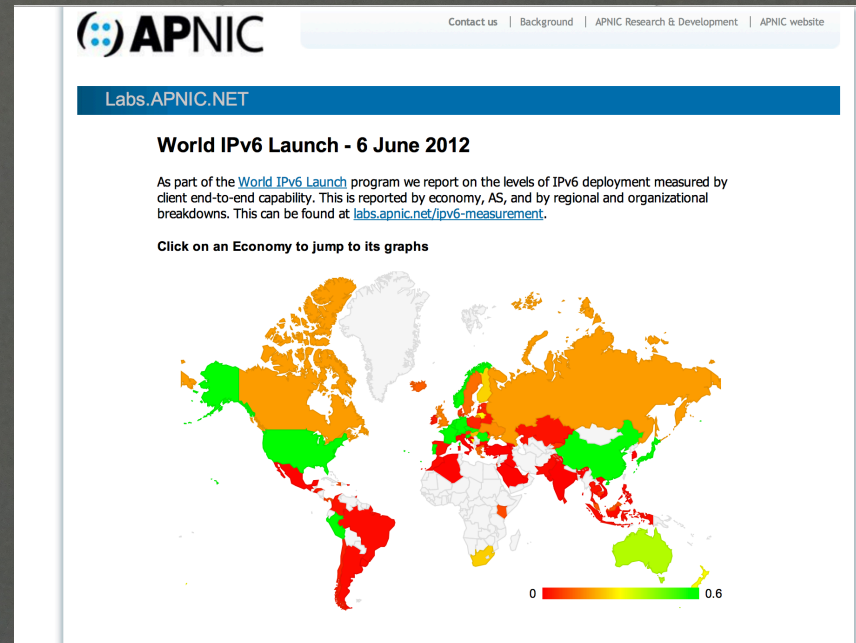


# Reports

We perform a basic scan of the daily data and produce a number of reports:

- a) A "summary" report of capabilities
- b) A map of the IPv6 world

<http://labs.apnic.net/index.shtml>





# Reports

We perform a basic scan of the daily data and produce a number of reports:

- a) A "summary" report of capabilities
- b) A map of the IPv6 world
- c) Per-ASN and Per-Country reports

<http://labs.apnic.net/ipv6-measurement/>



# Reports

We perform a basic scan of the daily data and produce a number of reports:

- a) A "summary" report of capabilities
- b) A map of the IPv6 world
- c) Per-ASN and Per-Country reports
- d) Daily Per-Country statistics report

<http://labs.apnic.net/dists/v6cc.html>



# The IPv6 Country League Table

Index	ISO-3166 Code	Internet Users	V6 Use ratio	V6 Users (Est)	Population	Country
1	RO	8655424	8.83%	764273	22080165	Romania
2	LU	469780	7.52%	35327	513983	Luxembourg
3	FR	50199417	5.57%	2796107	65025152	France
4	JP	100749839	4.02%	4050143	125937299	Japan
5	CH	6459171	3.05%	197004	7671225	Switzerland
6	DE	67931696	3.03%	2058330	82142318	Germany
7	US	249626028	2.60%	6490276	318807189	United States of America
8	BE	8504063	2.52%	214302	10447253	Belgium
9	CZ	7210339	1.80%	129786	10169732	Czech Republic
10	PE	10544079	1.74%	183466	30921055	Peru
11	NO	4588107	1.22%	55974	4720275	Norway
12	PT	5480254	0.82%	44938	10809180	Portugal
13	NL	15187783	0.77%	116945	16969591	Netherlands
14	TW	16213908	0.72%	116740	23162727	Taiwan
15	CN	566459146	0.63%	3568692	1348712254	China
16	SK	4348530	0.53%	23047	5490569	Slovakia
17	ZA	6816440	0.43%	29310	49039143	South Africa
18	SI	1416552	0.43%	6091	1995145	Slovenia
19	AU	19965255	0.40%	79861	22233024	Australia
20	SG	3718156	0.38%	14128	4816265	Singapore
21	GR	5057525	0.37%	18712	10783636	Greece
22	BA	1956050	0.34%	6650	4624232	Bosnia and Herzegovina
23	HK	4937241	0.30%	14811	7186669	Hong Kong Special Administrative Region of China
24	LT	2092635	0.26%	5440	3517034	Lithuania
25	GB	51951716	0.26%	135074	61773741	United Kingdom of Great Britain and Northern Ireland



# The IPv6 Country League Table

ISO-3166 Code	Internet Users	V6 Use ratio	V6 Users (Est) ▲	Population	Country
US	249626028	2.60%	6490276	318807189	United States of America
JP	100749839	4.02%	4050143	125937299	Japan
CN	566459146	0.63%	3568692	1348712254	China
FR	50199417	5.57%	2796107	65025152	France
DE	67931696	3.03%	2058330	82142318	Germany
RO	8655424	8.83%	764273	22080165	Romania
BE	8504063	2.52%	214302	10447253	Belgium
CH	6459171	3.05%	197004	7671225	Switzerland
PE	10544079	1.74%	183466	30921055	Peru
RU	60922920	0.24%	146215	137523522	Russian Federation
GB	51951716	0.26%	135074	61773741	United Kingdom of Great Britain and Northern Ireland
CZ	7210339	1.80%	129786	10169732	Czech Republic
NL	15187783	0.77%	116945	16969591	Netherlands
TW	16213908	0.72%	116740	23162727	Taiwan
AU	19965255	0.40%	79861	22233024	Australia
CA	28175891	0.25%	70439	34529279	Canada
NO	4588107	1.22%	55974	4720275	Norway
PT	5480254	0.82%	44938	10809180	Portugal
LU	469780	7.52%	35327	513983	Luxembourg
ID	56130719	0.06%	33678	250583569	Indonesia
ZA	6816440	0.43%	29310	49039143	South Africa
SK	4348530	0.53%	23047	5490569	Slovakia
GR	5057525	0.37%	18712	10783636	Greece
MY	16922715	0.11%	18614	27427416	Malaysia
UA	15123907	0.12%	18148	44613295	Ukraine



# The IPv6 ASN League Table

Economy	ASN	AS Name	# samples	v6 capable	v6 preferred ▼
<a href="#">US</a>	<a href="#">AS19782</a>	INDIANAGIGAPOP - Indiana University	1862	100	100
<a href="#">CN</a>	<a href="#">AS23910</a>	CNGI-CERNET2-AS-AP China Next Generation Internet CERNET2	1815	100	100
<a href="#">CN</a>	<a href="#">AS37944</a>	CNNIC-CSTNET-AP CHINA SCIENCE AND TECHNOLOGY NETWORK	354	100	100
<a href="#">AU</a>	<a href="#">AS38083</a>	CURTIN-UNI-AS-AP Curtin University	791	98.6094	97.9772
<a href="#">NZ</a>	<a href="#">AS24226</a>	CATALYST-IT-AS-AP Catalyst IT	925	95.7838	95.4595
<a href="#">US</a>	<a href="#">AS3598</a>	MICROSOFT-CORP-AS - Microsoft Corp	1234	74.5543	70.6645
<a href="#">JP</a>	<a href="#">AS55394</a>	GREE-NET GREE; Inc.	261	90.0383	66.6667
<a href="#">NZ</a>	<a href="#">AS58666</a>	NASL-AS-AP Network Access Services Limited	441	76.8707	66.2132
<a href="#">US</a>	<a href="#">AS1312</a>	VA-TECH-AS - Virginia Polytechnic Institute and State Univ.	588	69.2177	64.6259
<a href="#">NZ</a>	<a href="#">AS17649</a>	DMZGLOBAL-AP DMZGlobal Ltd	203	64.0394	64.0394
<a href="#">CZ</a>	<a href="#">AS197451</a>	VUTBR-AS Brno University of Technology	686	69.3878	63.9942
<a href="#">AU</a>	<a href="#">AS4608</a>	APNIC-AP Asia Pacific Network Information Centre	589	62.309	60.781
<a href="#">US</a>	<a href="#">AS6621</a>	HNS-DIRECPC - Hughes Network Systems	2001	60.4198	60.02
<a href="#">US</a>	<a href="#">AS91</a>	RPI-AS - Rensselaer Polytechnic Institute	236	61.4407	59.7458
<a href="#">RO</a>	<a href="#">AS12675</a>	UAIC-NETWORK Alexandru Ioan Cuza University	262	64.8855	57.2519
<a href="#">US</a>	<a href="#">AS5661</a>	USF - UNIVERSITY OF SOUTH FLORIDA	308	59.7403	56.4935
<a href="#">US</a>	<a href="#">AS2055</a>	LSU-1 - Louisiana State University	217	56.682	52.5346
<a href="#">GB</a>	<a href="#">AS786</a>	JANET The JNT Association	166359	60.7776	48.4242
<a href="#">BR</a>	<a href="#">AS22548</a>	N\xfacleo de Informa\xe7\xe3o e Coordena\xe7\xe3o do Ponto BR	302	58.2781	48.0132
<a href="#">NO</a>	<a href="#">AS57963</a>	LYNET-INTERNETT-AS Lynet Internett AS	304	53.6184	47.3684
<a href="#">US</a>	<a href="#">AS1351</a>	UVM-EDU-AS - University of Vermont	238	52.1008	47.0588
<a href="#">US</a>	<a href="#">AS6939</a>	HURRICANE - Hurricane Electric; Inc.	2152	43.9126	43.4015
<a href="#">US</a>	<a href="#">AS109</a>	CISCO-EU-109 Cisco Systems Global ASN - ARIN Assigned	569	57.4692	43.058
<a href="#">NZ</a>	<a href="#">AS18119</a>	ACSDATA-NZ ACSDData	983	46.4903	43.0315
<a href="#">US</a>	<a href="#">AS2698</a>	IASTATE-AS - Iowa State University	219	53.4247	42.9224
<a href="#">US</a>	<a href="#">AS6263</a>	NDIN - State of North Dakota	324	45.3704	42.284
<a href="#">HK</a>	<a href="#">AS4528</a>	HKU-AS-HK The University of Hong Kong	862	45.2436	41.8794
<a href="#">CN</a>	<a href="#">AS4538</a>	ERX-CERNET-BKB China Education and Research Network Center	5008	43.131	39.9561
<a href="#">SE</a>	<a href="#">AS12552</a>	IPO-EU IP-Only Telecommunication Networks AB	502	38.247	35.4582
<a href="#">CA</a>	<a href="#">AS16462</a>	UVIC-AS - University of Victoria	202	35.1485	35.1485



# And some Time Series...




Global IPv6



Europe IPv6




# And Some Countries...




Romania



France



Italy



Ireland



# What the...?

We noticed this class of entries in the web logs:

```
222.154.187.xx http://t10000.u1367873034830.s644708422.i647302.v10a.r6.td.labs.apnic.net/1x1.png
84.23.58.xx http://t10000.u1367873368824.s1566062113.i245974.v10i.r6.td.labs.apnic.net/1x1.png
```

We get some 200 of these web log entries every day

But \*.r6.td.labs.apnic.net has NO A record

So why are these clients attempting to fetch a V6-only URL using IPV4 as the transport protocol?

No idea!

No common country of origin, no common AS, no common device, no common host OS



# What the<sup>2</sup>...?

Even stranger...

```
202.124.201.xx http://[2401:2000:6660::f003]/1x1.png
118.148.0.xx http://[2401:2000:6660::f003]/1x1.png
```

We get some 16-20 of these web log entries every day

But this is a V6 literal form of URL!

Here's the origin Ases for this V4 fetch of a V6 literal URL for the 7th May

Origin AS	count	AS name
3352	1	TELEFONICA-DATA-ESPANA TELEFONICA DE ESPANA
4134	4	CHINANET-BACKBONE No.31,Jin-rong Street
4837	1	CHINA169-BACKBONE CNCGROUP China169 Backbone
35662	1	REDSTATION Redstation Limited
38793	8	NZCOMMS-AS-AP Two Degress Mobile Limited
55443	1	BAKST-AS-AU Level 16, 55 Hunter Street



# Questions?

APNIC Labs:

Geoff Huston

George Michaelson

[research@apnic.net](mailto:research@apnic.net)