

Measuring IPv6

Geoff Huston

APNIC Labs, August 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!

But...

- 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

Top videos

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

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

THE WEBBY AWARDS

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

travelberta.com

Find out more

Alberta
Canada

Ads are ubiquitous

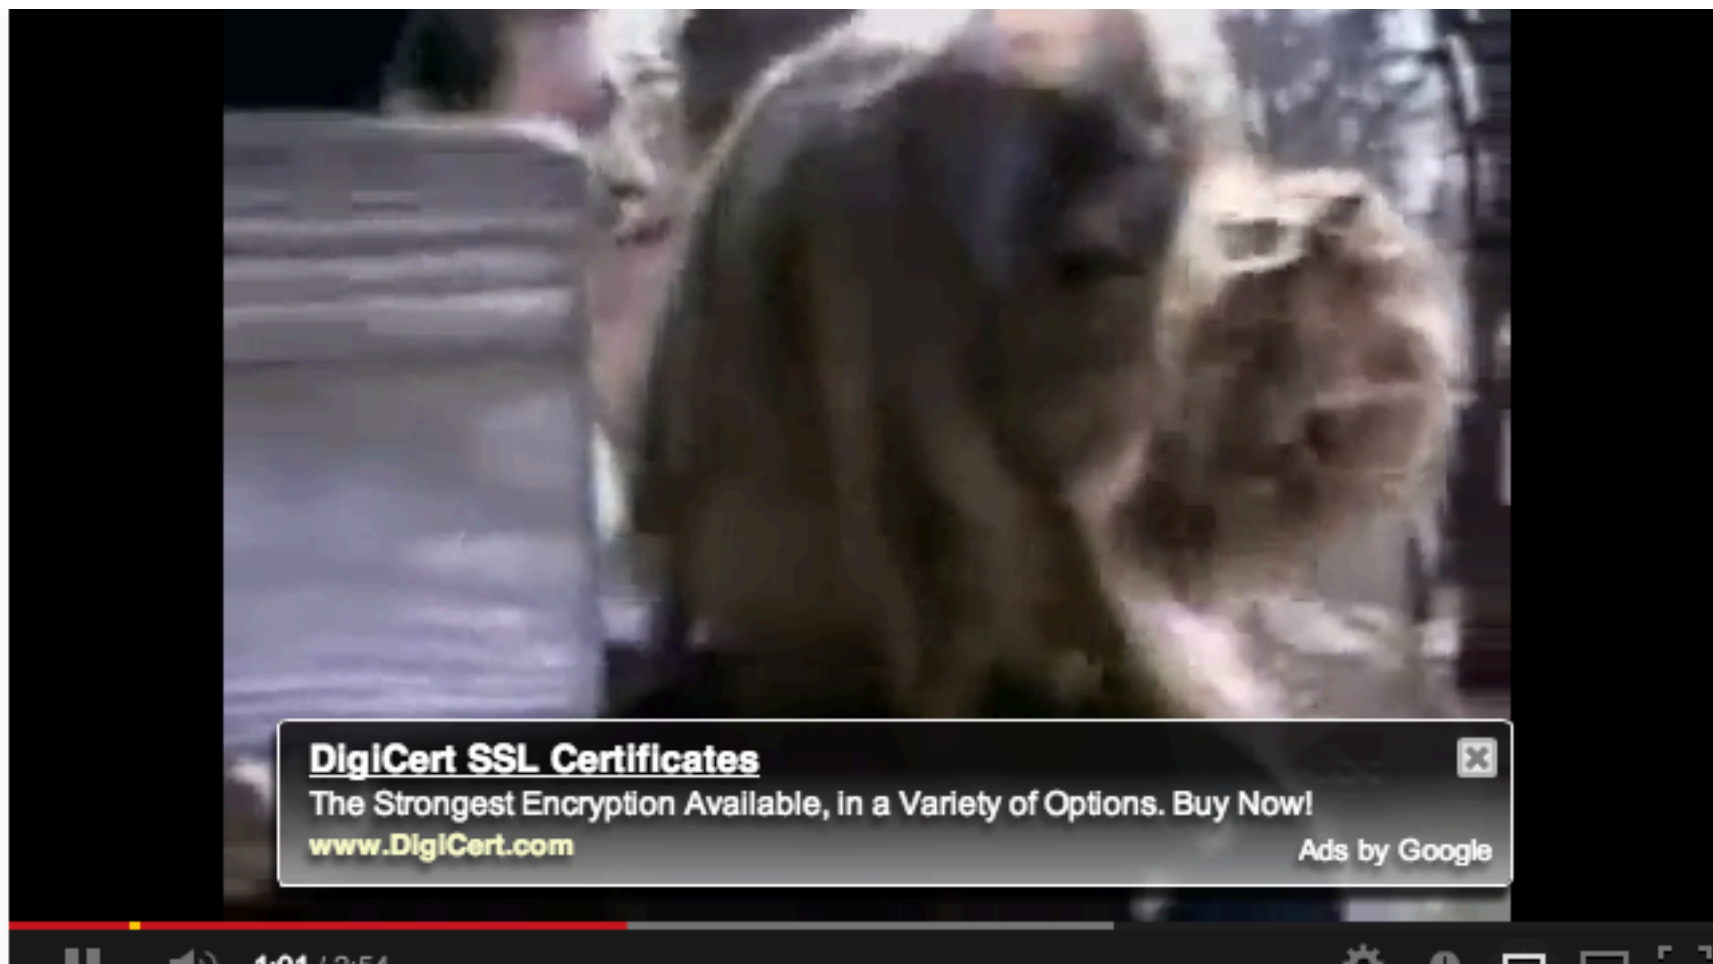
The screenshot shows a news website layout with several articles and advertisements. A large red circle highlights the following elements:

- Left sidebar:** A large advertisement for "REMINDER: SOMETIMES YOU NEED TO LET THE WILD OUT (remember to breathe)" featuring a landscape image of people on horseback.
- Top right:** A "travelalberta.com" advertisement with a "Find out more" button.
- Center:** A "W" logo advertisement for a "Low Rate Credit Card" with an "Apply now" button.
- Bottom right:** A "THE WOBBY AWARDS" advertisement.

Other visible content includes:

- Articles 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".
- Other advertisements like "Ballads of a thimble on", "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", and "AC Jimbo's European papers review".
- A "Top videos" section and a "Contact us" section.

Ads are ubiquitous



Ads are ubiquitous



Ads are implemented in Adobe Flash

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



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 '1x1' 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 click, 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
- Relatively 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.rand.apnic.net/measureipv6id.cgi?advertID=9999>)
 - 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

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/>



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>

Labs.APNIC.NET

World IPv6 Adoption

As a continuing activity following on from the [World IPv6 Launch](#) we report on the levels of IPv6 deployment measured by client end-to-end capability. This is reported by economy, AS, and by regional and organizational breakdowns. These can be found at labs.apnic.net/ipv6-measurement.

Click on an Economy to jump to its graphs



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	9715236	10.6587%	1035517	22070052	Romania
2	LU	468653	8.0921%	37923	515627	Luxembourg
3	EU	0	6.6177%	0	0	European Union
4	FR	51823998	5.7936%	3002475	65121888	France
5	CH	6539955	5.5336%	361894	7676004	Switzerland
6	JP	99484888	4.6031%	4579388	125850586	Japan
7	BE	8150710	4.2253%	344391	10449629	Belgium
8	DE	68166448	3.9990%	2725976	82128251	Germany
9	US	248890874	3.6797%	9158437	319664622	United States of America
10	PE	11168085	3.5454%	395953	31022460	Peru
11	SG	3427726	2.0627%	70703	4827784	Singapore
12	CZ	7418520	1.5100%	112019	10166535	Czech Republic
13	NL	15680146	0.8703%	136464	16988241	Netherlands
14	NO	4439774	0.8130%	36095	4724672	Norway
15	GR	5717112	0.7369%	42129	10787005	Greece
16	PT	6247622	0.6435%	40203	10816521	Portugal
17	SK	4088687	0.4651%	19016	5492595	Slovakia
18	UA	12785442	0.4525%	57854	44533065	Ukraine
19	AU	17732277	0.4480%	79440	22304751	Australia
20	NZ	3760114	0.3878%	14581	4372226	New Zealand
21	TW	16687082	0.3253%	54283	23176504	Taiwan
22	ZA	16656220	0.3201%	53316	49032147	South Africa
23	BA	2774697	0.3031%	8410	4624495	Bosnia and Herzegovina
24	KE	12123279	0.2964%	35933	43297426	Kenya
25	RU	67295058	0.2712%	182504	137336854	Russian Federation


The IPv6 Country League Table

Index	ISO-3166 Code	Internet Users	V6 Use ratio	V6 Users (Est) ▲	Population	Country
9	US	248890874	3.6797%	9158437	319664622	United States of America
6	JP	99484888	4.6031%	4579388	125850586	Japan
4	FR	51823998	5.7936%	3002475	65121888	France
8	DE	68166448	3.9990%	2725976	82128251	Germany
1	RO	9715236	10.6587%	1035517	22070052	Romania
30	CN	517268061	0.1625%	840560	1350569350	China
10	PE	11168085	3.5454%	395953	31022460	Peru
5	CH	6539955	5.5336%	361894	7676004	Switzerland
7	BE	8150710	4.2253%	344391	10449629	Belgium
25	RU	67295058	0.2712%	182504	137336854	Russian Federation
13	NL	15680146	0.8703%	136464	16988241	Netherlands
12	CZ	7418520	1.5100%	112019	10166535	Czech Republic
19	AU	17732277	0.4480%	79440	22304751	Australia
26	CA	28723173	0.2693%	77351	34606233	Canada
35	GB	53686776	0.1335%	71671	61822635	United Kingdom of Great Britain and Northern Ireland
11	SG	3427726	2.0627%	70703	4827784	Singapore
18	UA	12785442	0.4525%	57854	44533065	Ukraine
21	TW	16687082	0.3253%	54283	23176504	Taiwan
22	ZA	16656220	0.3201%	53316	49032147	South Africa
40	NG	45947414	0.1098%	50450	161615949	Nigeria
15	GR	5717112	0.7369%	42129	10787005	Greece
56	IN	123241492	0.0333%	41039	1223847993	India
16	PT	6247622	0.6435%	40203	10816521	Portugal
2	LU	468653	8.0921%	37923	515627	Luxembourg
14	NO	4439774	0.8130%	36095	4724672	Norway


The IPv6 ASN League Table

Economy	ASN	AS Name	# samples	v6 capable	v6 preferred
CN	AS23910	CNGI-CERNET2-AS-AP China Next Generation Internet CERNET2	347	100	100
US	AS19782	INDIANAGIGAPOP - Indiana University	533	99.8124	99.8124
AU	AS38083	CURTIN-UNI-AS-AP Curtin University	395	98.481	97.2152
AU	AS4608	APNIC-AP Asia Pacific Network Information Centre	347	96.83	91.0663
JP	AS55394	GREE-NET GREE; Inc.	261	89.6552	66.6667
US	AS15169	GOOGLE - Google Inc.	7371	80.057	7.5295
US	AS3598	MICROSOFT-CORP-AS - Microsoft Corp	594	76.936	69.8653
BR	AS22548	Nlxfacleo de Informalxe7ixe3o e Coordena\xe7\xe3o do Ponto BR	285	68.0702	62.1053
GB	AS786	JANET The JNT Association	84274	64.8812	55.1155
US	AS8071	MICROSOFT-CORP---MSN-AS-BLOCK - Microsoft Corp	355	63.3803	61.6901
US	AS6621	HNS-DIRECPC - Hughes Network Systems	1118	62.9696	66.458
US	AS109	CISCO-EU-109 Cisco Systems Global ASN - ARIN Assigned	323	58.2043	47.3684
AU	AS56132	MONASHUNI-AU-AS-AP Monash University;	889	50.3937	48.0315
CA	AS12093	UWATERLOO - University of Waterloo	209	42.5837	40.6699
HK	AS4528	HKU-AS-HK The University of Hong Kong	258	40.3101	37.9845
US	AS87	INDIANA-AS - Indiana University	1302	39.9386	36.1751
CN	AS7497	CSTNET-AS-AP Computer Network Information Center	458	38.4279	31.8777
SG	AS24482	SGGS-AS-AP SG.GS	266	37.594	37.594
US	AS22394	CELLCO - Cellco Partnership DBA Verizon Wireless	11565	35.668	33.1431
BE	AS12392	ASBRUTELE Brutele SC	5448	34.8201	33.2048
JP	AS2516	KDDI KDDI CORPORATION	35190	31.2191	29.1162
SE	AS12552	IPO-EU IP-Only Telecommunication Networks AB	272	30.5147	30.5147
MY	AS17564	GITN-PCN-AS-AP GITN (M) Sdn. Bhd.	236	30.0847	25.8475
SG	AS7472	NUS-AS-AP Computer Centre	216	29.1667	23.1481
JP	AS18126	CTCX Chubu Telecommunications Company; Inc.	3380	29.0533	26.8343
DE	AS12816	MWN-AS Leibniz-Rechenzentrum	313	28.4345	24.6006
CN	AS4538	ERX-CERNET-BKB China Education and Research Network Center	1520	27.5658	25.1316
GR	AS3323	NTUA National Technical University of Athens	631	27.4168	26.6244

And some Time Series...




Global IPv6



Asia IPv6


And Some Countries...

IPv6 Preference by Month




Romania

IPv6 Preference by Month




USA

IPv6 Preference by Month



China

IPv6 Preference by Month



Japan

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!

What the²....?

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 Degrass Mobile Limited
55443	1	BAKST-AS-AU Level 16, 55 Hunter Street

Questions?

APNIC Labs:

Geoff Huston

George Michaelson

research@apnic.net