

Measuring IPv6 with advertisements for fun and profit

George Michaelson, APNIC

how to measure “the Internet”

- What do we mean when we say “we’re measuring the Internet?”
 - Number of hosts
 - Number of routes
 - Number of active routing entities
 - Number of application-events
 - Number of voip calls
 - Number of voice-over-IP in the carrier calls
 - Number of streaming TV watchers
 -

Clearly a multi-dimensional problem

- We're well beyond single measure
 - Routing scaling for routing professionals
 - Traffic volumes for peering/settlement
 - Voice/data mix for telephony nuts
- Finding datasets to suit the context
 - DiTL has root-DNS traffic covered
 - IX captures get local traffic mix, indicative volumes
 - ISPs track their own customer
 - Websites track their own users
- But this appears too valuable to share widely

Who is measuring the end user?

Measuring the end user for IPv6 readiness

- Need a technique which is ubiquitous
- “Hop over” missing IPv6 enabled CPE
- Reflect real-world end-user behaviour
- Avoid systemic bias ‘my own web is typical’
 - Demonstrably not true for technical community
 - Access by tech end-user is ‘near the core’ bypassing CPE
 - Often has ‘special’ access (new/novel/experimental & pre-service release)

Javascript

- Small community of researchers using javascript to ‘tickle’ web users via undisplayed fetches
 - Additions to website, not rendered into DOM, using unique DNS names, trackable
 - *.test.domain wildcards sub-classed to provide unique per-client test names
 - Javascript driven client-side view of delay
 - Summary sent to web by ‘fetch’ of DNS name
 - Data in the name fetched, not in the content sent
- Not substantially different to normal website tracking methodologies
 - Third party site, or alternate DNS names to main site

Anatomy of a javascript measurement

- Website markup includes .js fetch, or inline `<script>...</script>` block
 - Javascript engine in browser runs asynchronously to page render
 - Web page drawn unaffected/in-parallel with fetches
 - Spin random, to indentify test. All fetches include unique id in wildcard DNS space
 - test314159.<test>.labs.apnic.net
 - Series of test images fetched in sequence (or random)
 - Dual-stack test314159.rdtd.labs.apnic.net
 - IPv6 only test314159.rdt6.labs.apnic.net
 - IPv6 literal test314159.v6lit.labs.apnic.net
 - Each fetch has its own 'sprite' like timer
 - On completion, client-side delay (ms) measured from base clock
 - Fall timer, to send results at hang time if tests don't complete (10s)
 - Results returned with same unique test id
 - test314159.zrdtd44.zrdtd6101.zv6litnull.results.labs.apnic.net

What do we get

- Configure DNS to have single NS, host NS and turn on query logging
 - DNS logs, client fetching DNS names to begin tests with timers
- TCPdump of packetflows to webserver, dns, tunnel endpoint
 - Can detect partial connect failures, ICMP/ICMPv6 & SYN flows, TCP mss.
 - Also detailed inter-packet timings
- Web logs
 - Successful fetch logging, order not guaranteed in logfile but has server-side timing information

Cross-collating the data

- Initial .js fetch always on IPv4.
 - Confirms IPv4 address under test
- Final results.labs.apnic.net posting also always on IPv4
 - Confirms test sequence ran to completion
 - If received, also has client-side delay times
- All tests from same host carry same random ID
 - Where logged, can then cross-correlate IPv4 and IPv6
- IPv4 and IPv6 can be seen in TCPdumps
- Cross index to DNS resolver IP in DNS query also possible

Post-test processing

- results line received before tests complete
- results line never received, but individual tests run to completion
- tests lag by extreme periods, minutes after test
 - (so contradict results line which says null for that test)
- Post-process heuristics manage this to produce ‘unified’ log of test combining data from web log and TCPdump
 - If any source said we saw the test, its included, even if results say test wasn’t run (!)
 - If results provide times, then these times are used, otherwise server-side times are used.

What do we get?

- Outcome: measurements at the 50,000 -100,000 hits/day level across 20-30 participating websites
 - large hits from specific economies/websites, skewing data capture
- Still valid, but not yet ‘global’
 - A site like wikipedia, or an international newspaper would be **EXTREMELY INTERESTING** as a collection source
 - Javascript can be used to perform 1-in-1000 type sub-rate filters to sample highly popular sites

Solutions

- Need a way to get global coverage
- Want to leverage javascript investment, use same data collection methodology (combine datasets)
- Looking for a vehicle similar to javascript, but not limited to websites we can persuade to include our code.

...buy the impressions

- Web advertising networks now fundamental to 'making the web pay'
 - *Lots of websites willing to have adverts placed for \$*
- Well designed framework for distribution of content to websites en masse
 - *Submit once, placement as widely as possible **worldwide***
- Simple payment model based on impressions/clicks
pricepoint: CPM (clicks per mille)
 - *Low CPM translates to high impression count*
 - Remember: the advertising network wants your money, so if you bid too low for clicks, you get sold placements, to justify the payment
 - Daily investment of \$20 buys 50,000 impressions/day

From javascript to flash

- Advertising using flash encoded 'dynamic' content
 - Flash authoring tools widely available, ubiquitous for dynamic website content
 - Advertisement presents as an 'image' but can use flash to download active elements, movie clips, &c
- The advertising economy is now almost completely based on flash
 - Turn off flash, or run an adblocker and visit your normal daily diet of websites...
- Javascript and actionscript near-cousins
 - Simple to translate working javascript into flash

Minor Problems

- Advertising networks limit which flash primitives you can use
 - Luckily, `fetchURL()` is basic
 - Exclude `random()` library calls, but provide information into flash advert which is functionally highly random, can reduce via `crc32()` type hash
 - Cross site scripting attacks demand use of `crossdomain.xml` fetch
 - Could almost base method on this fetch alone (!)
 - Twice the volume of fetches required for same measurement
 - Flash not on all platforms (cannot measure iOS)

Placement


- At low CPM, advertising network needs to present unique, new eyeballs to harvest impressions
 - Therefore, ‘good’ advertising network provides fresh crop of unique clients per day
- Language-specific selections can tune placement
 - Evidence suggests that of 250 iso3166 economies, we have secured placement into 200, with 150+ at significant volume

Unique IPS?

- Collect list of unique IP addresses seen
 - Per day
 - Since inception
- Plot to see behaviours of system
 - Do we see 'same eyeballs' all the time?


Lots of Unique IP'S

google uniques/day + google cumulative uniques * javascript uniques/day * javascript cumulative unique -




Lots of Unique IP'S

google uniques/day + google cumulative uniques * javascript uniques/day * javascript cumulative unique -



Lots of Unique IP'S

google uniques/day + google cumulative uniques * javascript uniques/day * javascript cumulative unique -



Lots of Unique IP'S

- Both Javascript and Flash delivering consistent, variant IP address sources every day
- Slight sign of bias in Javascript
- Google Ads placement demands 'fresh eyeballs' to justify charge of placement based on impressions
 - They do a good job: how do you think they know the eyeballs haven't been seen before?
 -

Lots of Unique IP'S


- Both Javascript and Flash delivering consistent, variant IP address sources every day
- Slight sign of bias in Javascript
- Google Ads placement demands 'fresh eyeballs' to justify charge of placement based on impressions
 - They do a good job: how do you think they know the eyeballs haven't been seen before?
 - (they ask google!)

ASN Coverage

- Collect list of unique AS for both IPv4 and IPv6 seen each day
- Collate since Inception
 - Plot to see behaviours


IPv4 ASN Coverage by time

google uniques/day + javascript uniques/day combined cumulative/day
google cumulative uniques * javascript cumulative unique combined cumulative unique



IPv6 ASN Coverage by time

combined cumulative/day + combined cumulative unique *



AS Range by time

- Javascript shows more signs of ‘weekend droop’
 - Google aim to supply consistent load over time so artificially ‘inflate’ traffic against normal usage
 - Remember we’re $1/n$ th for a very small n of total advertising, so they can make us ‘constant’ when the javascript reflects real-world load per website and so can’t mask the ‘drop-off’ of weekend load
- Trending to 25,000+ AS seen in IPv4
- vs 1400 in IPv6
 - Few AS unique to either collection method (js/flash)

AS Range is Representative

- 25,000 ASN in IPv4 is a SIGNIFICANT amount of the global DFZ routespace
- We believe there is considerable headroom in the unique IP served by the advertizing network
- We believe that we can use this data to make observations about global/internet-wide behaviours, at the end-user.
 - Worldwide

Dealing with the data

- Per-day, unified web, dns, tcp dumps
- Single-line per IPv4/IPv6 instance (client being tested)
 - Times of dual-stack, IPv6 literal, relative to IPv4 fetches
 - Approx 5Mb per day, 250,000 experiments/day
- Post-process to add
 - Economy of registration (RIR delegated stats)
 - Covering prefix and origin-AS (bgp logs for that day)
- Combine into weekly, monthly datasets (<5Mb)


Wadja finding?

- labs.apnic.net/ipv6_measurement
 - Breakdowns by ASN, Economy, Region, Organisation
 - JSON and CSV datasets for every graph on stable URL
- 129 economies provide >200 samples/interval consistently in weeklies, 153 at monthlies.
 - Law of diminishing returns as more data collected
 - 200 is somewhat arbitrary, but provides for 0.005 level measure if we get one-in-200 hit. Beyond this, its not easy to justify data is sufficient to measure lowside
- Lets go look at some of the data...

World trend

001 World (935891 samples/week) +

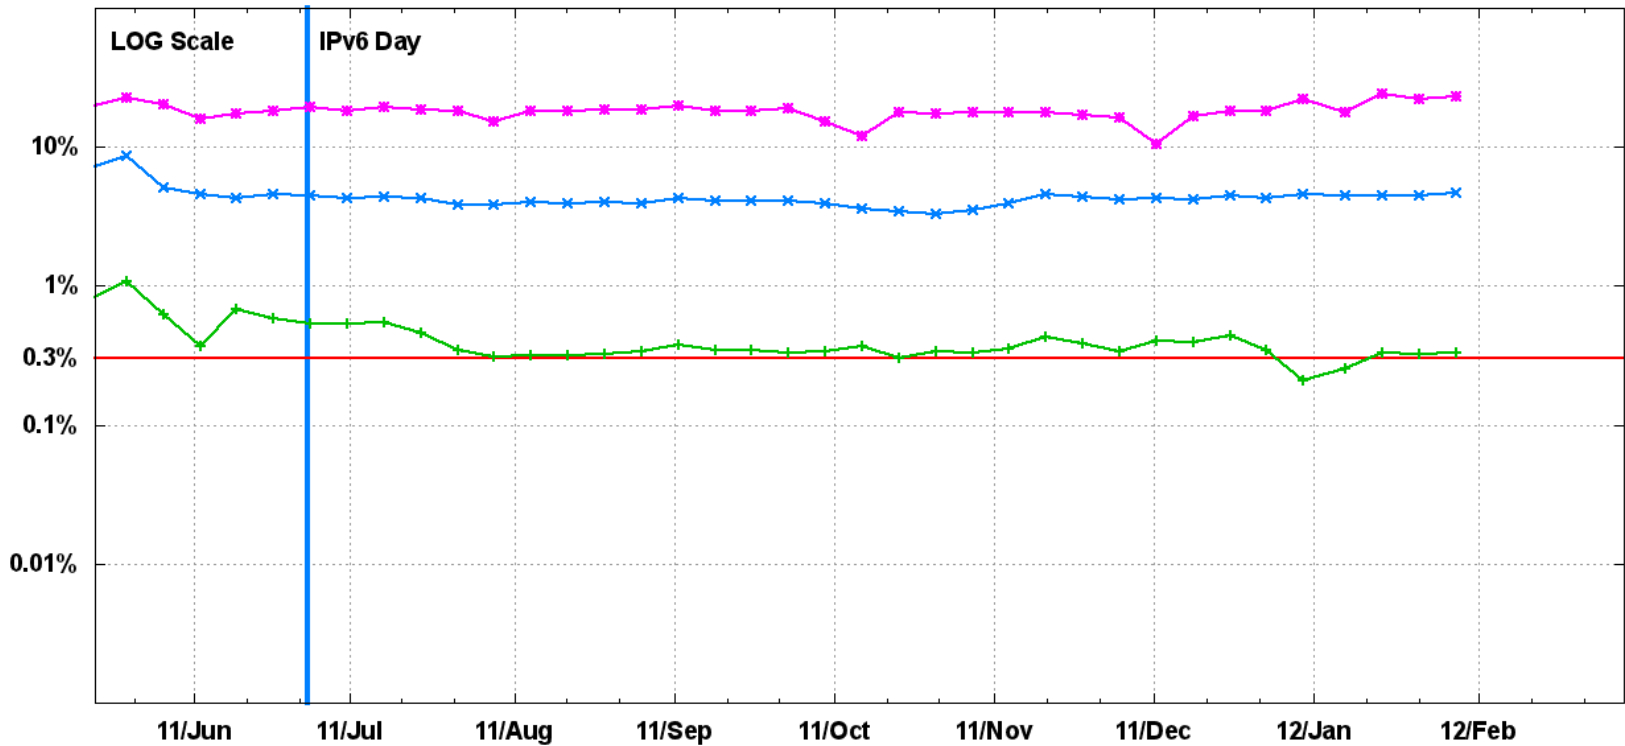
001 World IPv6 preferred (as % v4)



World trend

IPv6 preferred + IPv6 capable * IPv6 coerceable *

001 World IPv6 preferred (as % v4) from 935891 samples/week




The world is measured as 0.3%

- 900,000 samples per week informing a 0.3% IPv6 preference level
- Minor variations, no obvious trend upward
- Co-ercable/Capable figures track Vista/W7 penetration into the desktop, support 40% capability figure, if IPv6 CPE deployed
- Good match to other figures for worldwide trend

Drill down Europe

150 Europe (225159 samples/week) +

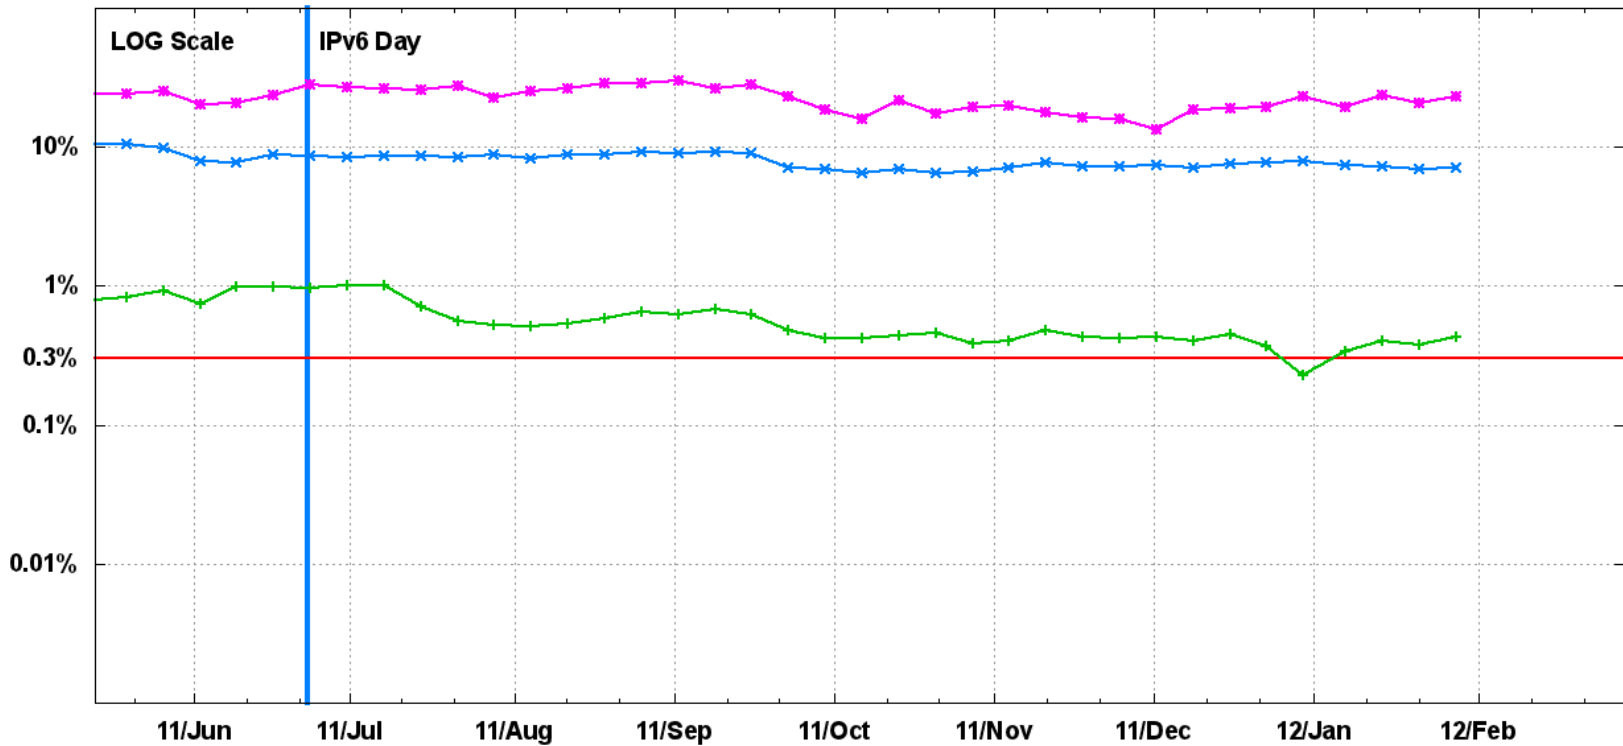
150 Europe IPv6 preferred (as % v4)



Drill down Europe

IPv6 preferred + IPv6 capable * IPv6 coerceable *


150 Europe IPv6 preferred (as % v4) from 225159 samples/week



Drill down Western Europe

155 Western Europe (20069 samples/week) +


155 Western Europe IPv6 preferred (as % v4)



Drill down Western Europe

IPv6 preferred + IPv6 capable * IPv6 coerceable *

155 Western Europe IPv6 preferred (as % v4) from 20069 samples/week



Observations

- The world is 'lumpy' for IPv6 preference
 - We can detect regional/economy-specific variances against world 0.3%
- Sample sizes for OECD economies, UN regions are good enough to continue
 - Increased uptake would help extend coverage for iso3166, observing we're already at 150/250 for some data, and over 100 for 'good' data volumes

AS views

- Take per-day BGP views (AS4608)
- Process IP stream by longest-match covering prefix, emit prefix and Origin AS
- Result: per Origin-AS views of IPv6 preference
 - 1470 pass the 200-minimum samples test
 - Out of a population of 5000 ASN in the IPv6 DFZ
 - more data required to increase AS specific measure, but we're probably not at the law of diminishing returns yet: I believe double the data would increase this markedly

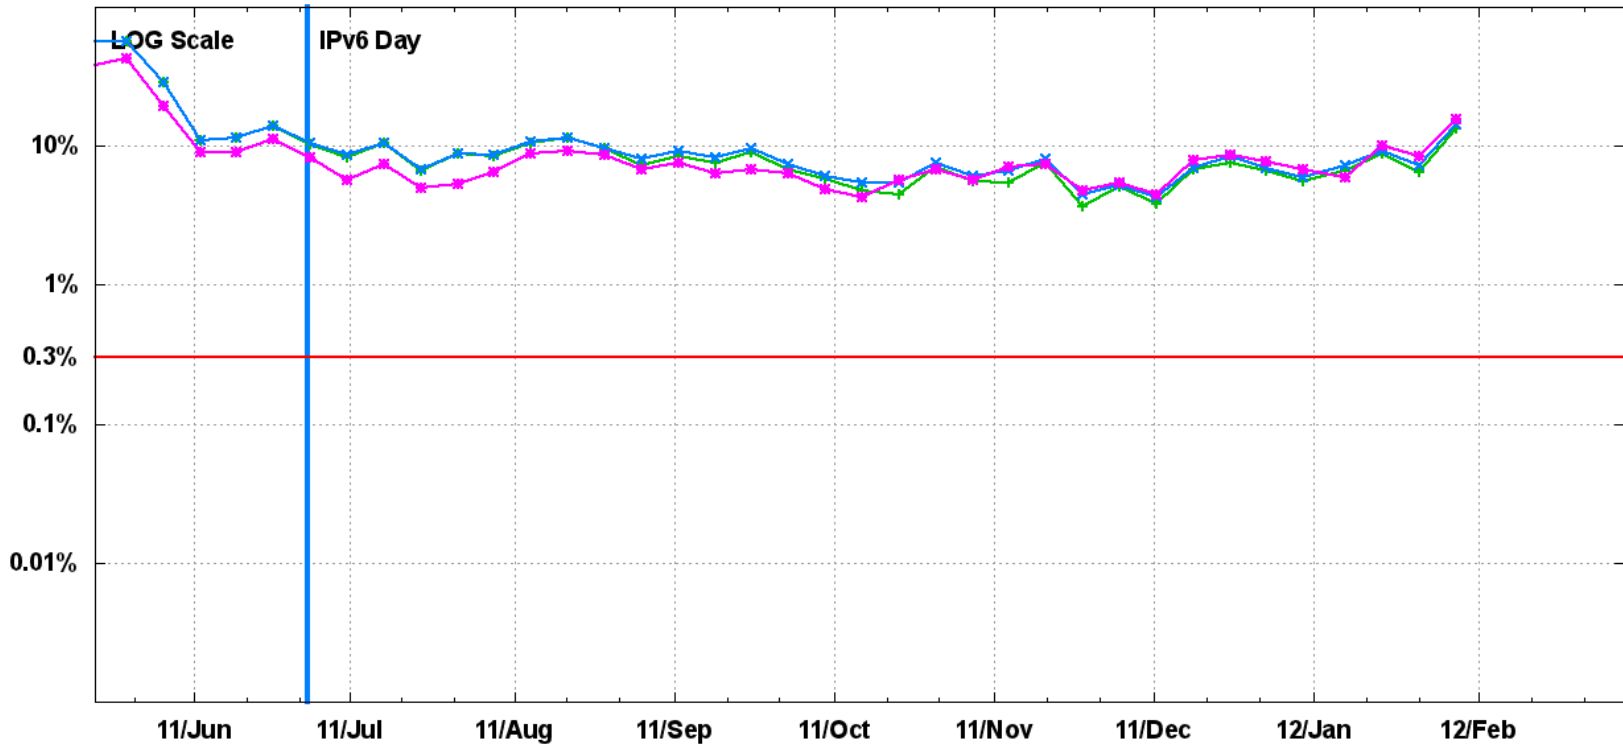
Two kinds of IPv6 AS

- Initial data suggests two kinds of AS visible in the IPv6 measure
- ASN which have true IPv6 delivered to the end user
 - Preferred, capable, coercible data co-incident
- ASN with the CPE roadblock
 - Divergeant preferred compared to capable/coercible

Hurricane Electric (6939)

IPv6 preferred + IPv6 capable * IPv6 coerceable *

AS6939 IPv6 preferred (as % v4) from 981 samples/week




Hurricane Electric (6939)

- Few direct customers/CPE issues
- Many tunnel endpoints homed in HE
- Many peerings, IPv6 presented at many IX
- Transit role

Free SAS / Proxad (12322)

AS12322 (1482 samples/week) +


AS12322 IPv6 preferred (as % v4)



Free SAS / Proxad (12322)

IPv6 preferred + IPv6 capable * IPv6 coerceable *

AS12322 IPv6 preferred (as % v4) from 1482 samples/week




Free SAS / Proxad (12322)

- 6RD deployed to customers
- High penetration of IPv6 Preference
 - But distinctly less than Capable/Coercible
 - reflects the Vista/W7 population inside the ISP who haven't taken up a 6RD option

Comcast (7922)

AS7922 (10555 samples/week) —+


AS7922 IPv6 preferred (as % v4)



Comcast (7922)

IPv6 preferred + IPv6 capable * IPv6 coerceable *

AS7922 IPv6 preferred (as % v4) from 10555 samples/week



Comcast (7922)

- Comcast using several ASN, not all seen in this measurement
 - Some not used for customer facing nets
- IPv6 deployment is in early stages, most customers would still present behind IPv4 only CPE, hence the gap between preference and capable/coerceable.

Finding the hopovers

- Find list of IPv6 prefixes we see in test with their respective IPv6 Origin-AS
 - Compare to the matching IPv4 they present from, for its Origin-AS
 - Either the same AS, or different.
- Divergent may mean the ISP has used another provider as a ‘hop-over’
 - because they don’t have native IPv6 transit, or end-user experimentation or ... ?

Top V6 'hopped into' ASN

ASN	Nice	# used IPv6 ASN	Relative weight
9264	ASNET Academic Sinica Network	31	0.017
9264	UNINET-TH	34	0.019
1659	ERX-TANET-ASN1 Tiawan Academic Network (TANet) Information Center	40	0.0223
7018	ATT-INTERNET4 - AT&T Services; Inc	46	0.026
12989	HWNG Eweka Internet Services B.V	51	0.028
23910	CNGI-CERNET2-AS-AP China Next Generation Internet CERNET2	51	0.028
2852	CESNET2 CESNET; z.s.p.o.	52	0.029
2516	KDDI KDDI CORPORATION	148	0.083
6939	HURRICANE - Hurricane Electric; Inc.	308	0.172
12322	PROXAD Free SAS	316	0.176

Observations

- Proxad/Hurricane offering tunnel broker, not ad-hoc (teredo/6to4) off their own IPv6 prefix
- Mobile users, on the road, but with static tunnel definition
- High count from China ...
 - High count inside china too.
 - Far higher.
- Where are the IPv4 nets they are hopping 'in' from?

Whats going on in China?

Economy	#	Economy	#
BD	1	MY	6
CA	1	TH	6
CL	1	JP	7
GR	1	EU	8
IN	1	ID	11
PK	2	US	16
SG	2	TW	21
AU	3	VN	48
KR	3	CN	443
PH	4		

Top V4 'hopped from' ASN

ASN	Nice	# used IPv6 ASN	Relative weight
3320	DTAG Deutsche Telekom	14	0.013
4837	CHINA169-BACKBONE CNCGROUP China169 Backbone	14	0.013
6799	OTENET-GR Ote SA (Hellenic Telecommunications Organisation)	14	0.013
7738	Telecomunicacoes da Bahia S.A.	14	0.013
6830	UPC UPC Broadband	15	0.014
28573	NET Servicos de Comunicacao S.A.	16	0.015
4713	OCN NTT Communications Corperation	17	0.016
3269	ASN-IBSNAZ Telecom Italia S.p.a.	19	0.018
4134	CHINANET-BACKBONE No.31;Jin-rong Street	26	0.024
8151	Uninet S.A. de C.V.	36	0.037

Observations

- Again, promiscuous hop-over tunnels using a range of endpoint IPv6 transit/tunnels.
- Not Teredo/6to4, so hand-installed or private arrangements
- We know some of these people have IPv6 native..

Conclusions

- JS on web, and Flash in advertising networks viable for broad-range high volume data collection
- Low cost of entry, high return on investment for measurement
- Internet wide, unique IP Addresses. 20-30,000 ASN (v4) to 2000 ASN (v6) visible
 - Good, 50%+ coverage of V4, 30%+ of IPv6 active ASN
- We have a LOT more information to get out of this investment. RTT, MTU/MSS, pMTU &c &c

IPv6 measurement

- Penetration rate of IPv6 into the global AS economy is slow
- No signs of ‘game changer’ behaviour at the end-user yet
- Widely distributed hop-over behaviour emerging, due to the CPE gap ?
- Even IPv6 enabled ISPs have people seeking to tunnel over the air-gap
- Much more information about IPv6 in the data
 - “watch this space” –longterm investment in measurement, ongoing.

Acknowledgements

Thanks to RIPE NCC, Google and ISC

Acknowledgements

Thanks to RIPE NCC, Google and ISC

More info, JSON/CSV data (daily updates)

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

?

