Operating the Internet's Largest Measurement System

Geoff Huston
APNIC Labs

Indian ISPs

Visible ASNs: Customer Populations (Est.)

Date: 07/12/2024

Rank	ASN	AS Name	CC	Users (est.)	% of country	% of Internet	Samples
1	AS55836	RELIANCEJIO-IN Reliance Jio Infocomm Limited	IN	291,685,503	48.71	6.905	140,538,196
2	AS45609	BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service	IN	147,672,751	24.66	3.496	71,150,818
3	AS24560	AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services	IN	30,568,128	5.1	0.724	14,728,156
4	AS38266	VIL-AS-AP Vodafone Idea Ltd	IN	27,143,865	4.53	0.643	13,078,298
5	AS45271	ICLNET-AS-AP Idea Cellular Limited	IN	17,878,027	2.99	0.423	8,613,886
6	AS9829	BSNL-NIB National Internet Backbone	IN	11,785,640	1.97	0.279	5,678,488
7	AS24309	CABLELITE-AS-AP Atria Convergence Technologies Pvt. Ltd. Broadband Internet Service Provider INDIA	IN	4,063,322	0.68	0.096	1,957,766
8	AS133982	EXCITEL-AS-IN Excitel Broadband Private Limited	IN	3,599,305	0.6	0.085	1,734,196
9	AS45916	GTPL-AS-AP Gujarat Telelink Pvt Ltd	IN	3,403,569	0.57	0.081	1,639,888
10	AS133661	NETPLUS-AS Netplus Broadband Services Private Limited	IN	3,186,596	0.53	0.075	1,535,347

Use of IPv6 in India

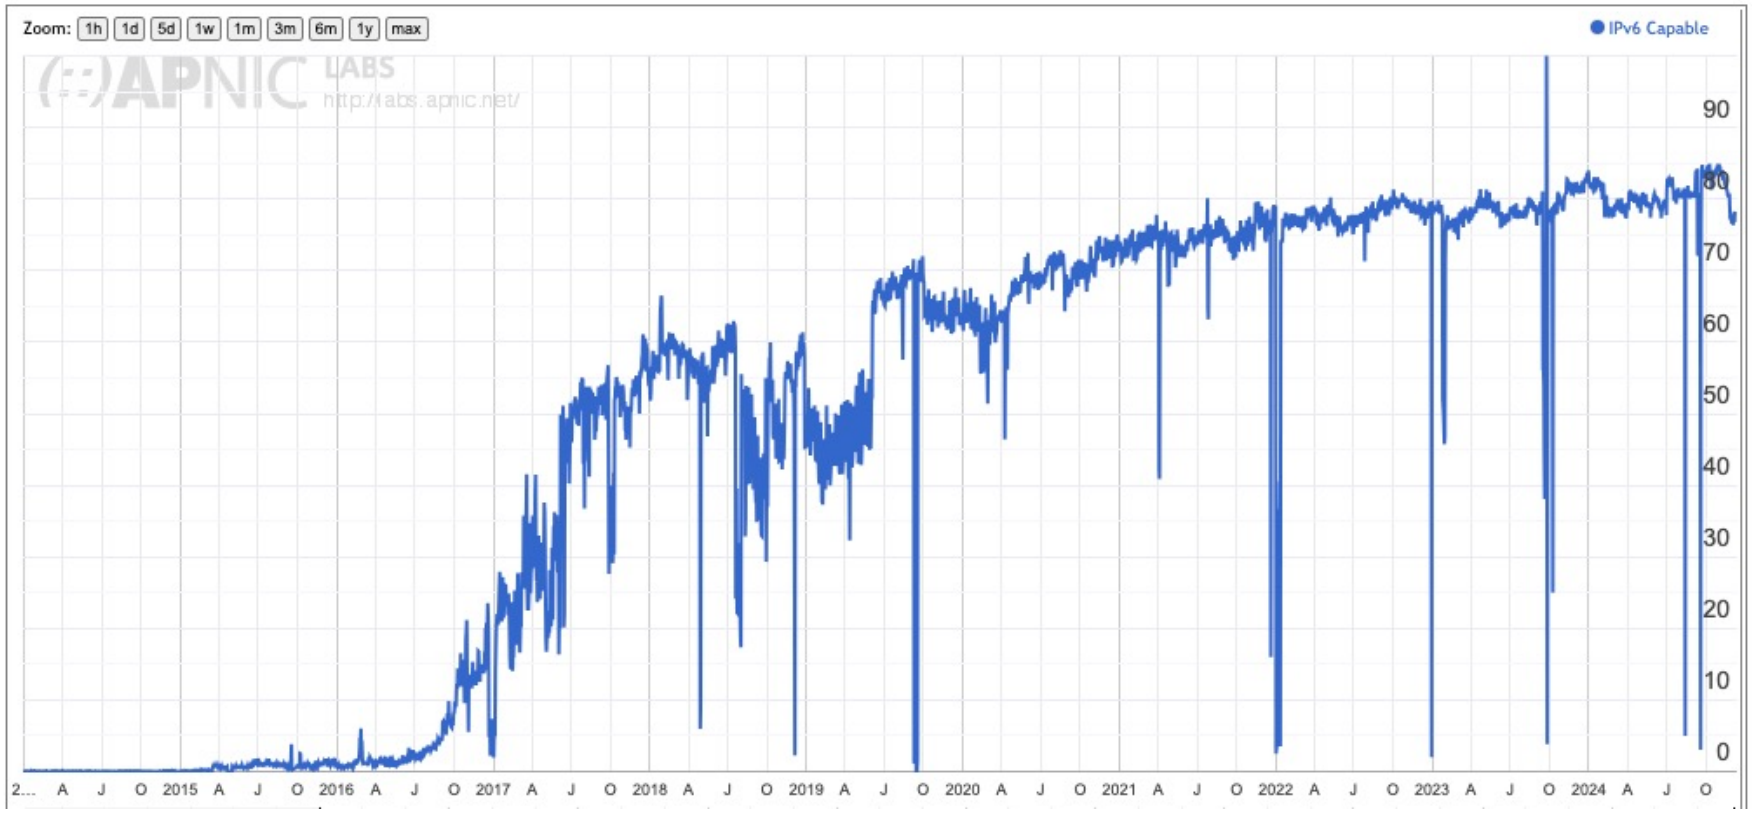

Use of DNSSEC in India

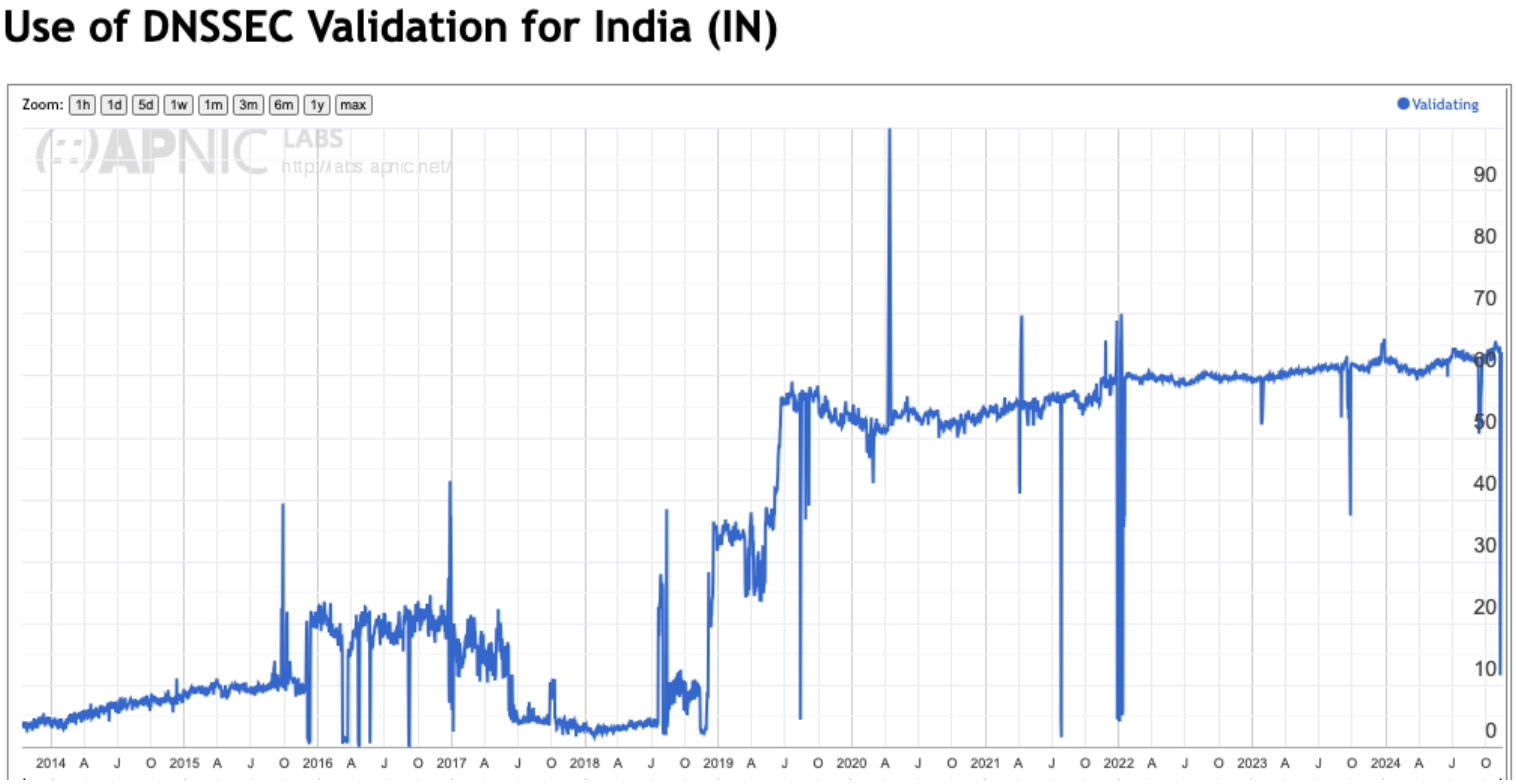

Use of RPKI in India

Use of Route Object Validation for India (IN)

Display: **Addresses** (Advertised ROA-Valid Advertised Addresses), **IPv4**, **Percent** (of Total)

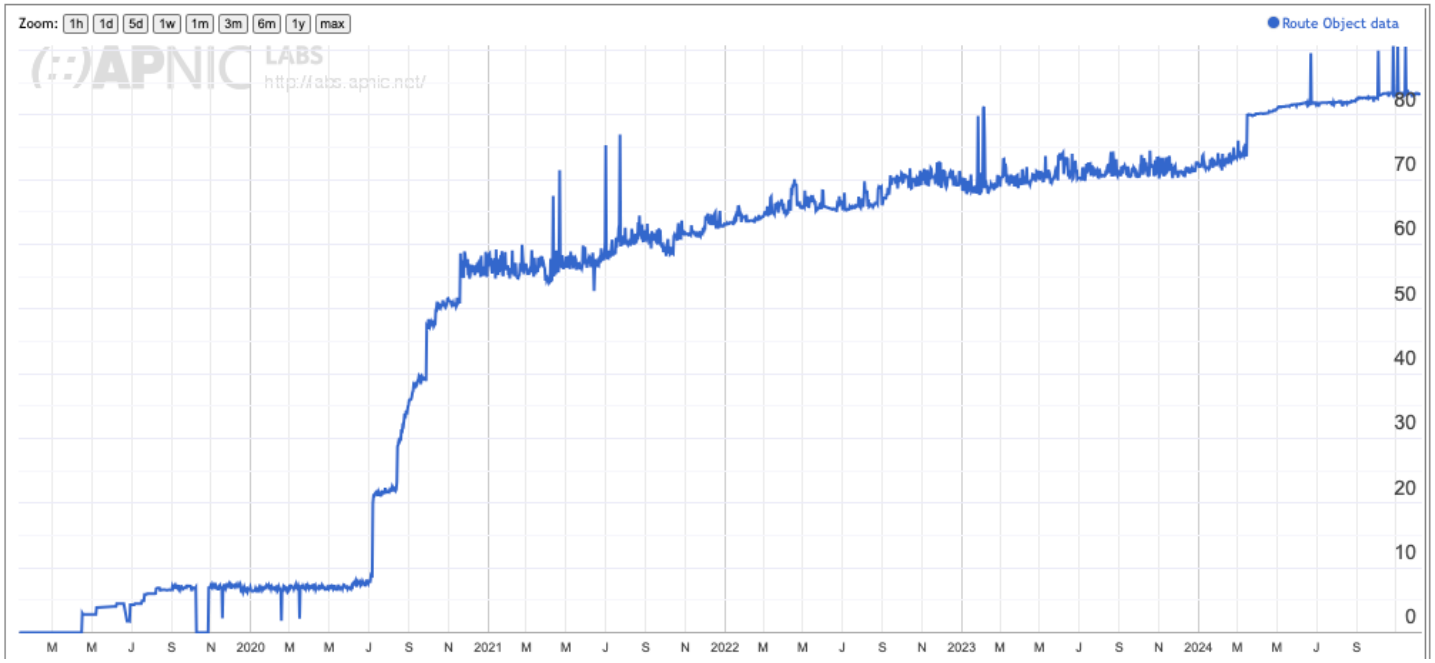

Use of ROV-Drop in India

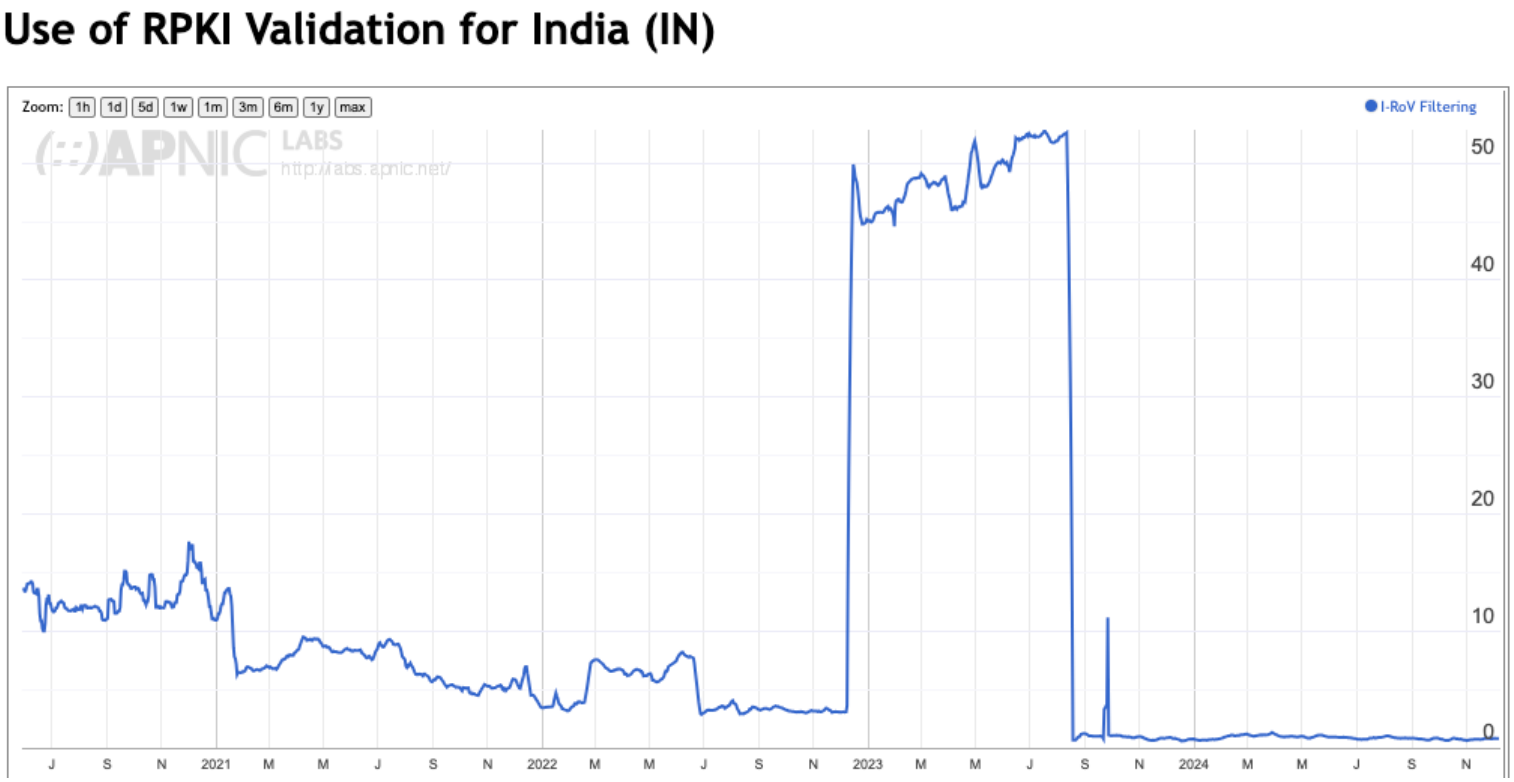

So much to measure…

so little time!

How and what

- The basic questions are to figure out **what** measurement approach to use:
  - whole of network metrics vs sample measurements
  - system behaviour metrics vs user experience metrics vs component metrics
  - snapshot metrics vs time series measurements

How and what

- The basic questions are to figure out **what** measurement approach to use:
  - whole of network metrics vs sample measurements
  - system behaviour metrics vs user experience metrics vs component metrics
  - snapshot metrics vs time series measurements

- And then figure out **how** to perform the measurement:
  - Passive vs active
  - Dedicated probes vs enrolment
  - Deliberate and Intentional vs Opportunistic

APNIC's path into measurement

2010:

IPv4 exhaustion is just around the corner - so how “ready” are we to operate an IPv6 Internet?

"Measurable" IPv6 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?

…

Close, but

- None of these specific measurement questions really embrace the larger question of IPv6 “readiness” for the Internet as a whole
  - 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

What's the 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?

- What if we use scripting on a web server to text the capabilities of clients via a scripted set of related web object fetches
  - That way we can test a very large number of clients for IPv6

Scale of 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
  - They tend to be extremely suspicious of adding script elements to their service that performs third party dual stack tests with their clients (and I can’t blame them!)

Scale of Measurement

- We really need to use a massively [illegible] web service to conduct this experiment
  - But "massively popul[illegible]rvices" worry constantly about service resiliency and p[illegible]their data regarding users
  - They ter[illegible]extremely suspicious of adding script elements to their se[illegible]performs third party dual stack tests with their clients (and I can't bl[illegible]them!)

We need to rethink this approach…

How to conduct measurements at scale

Be Google

How to conduct measurements at scale

Or get Google to place your measurement code on millions of end user’s systems, all of the time

How to conduct measurements at scale

How?

Online Ads

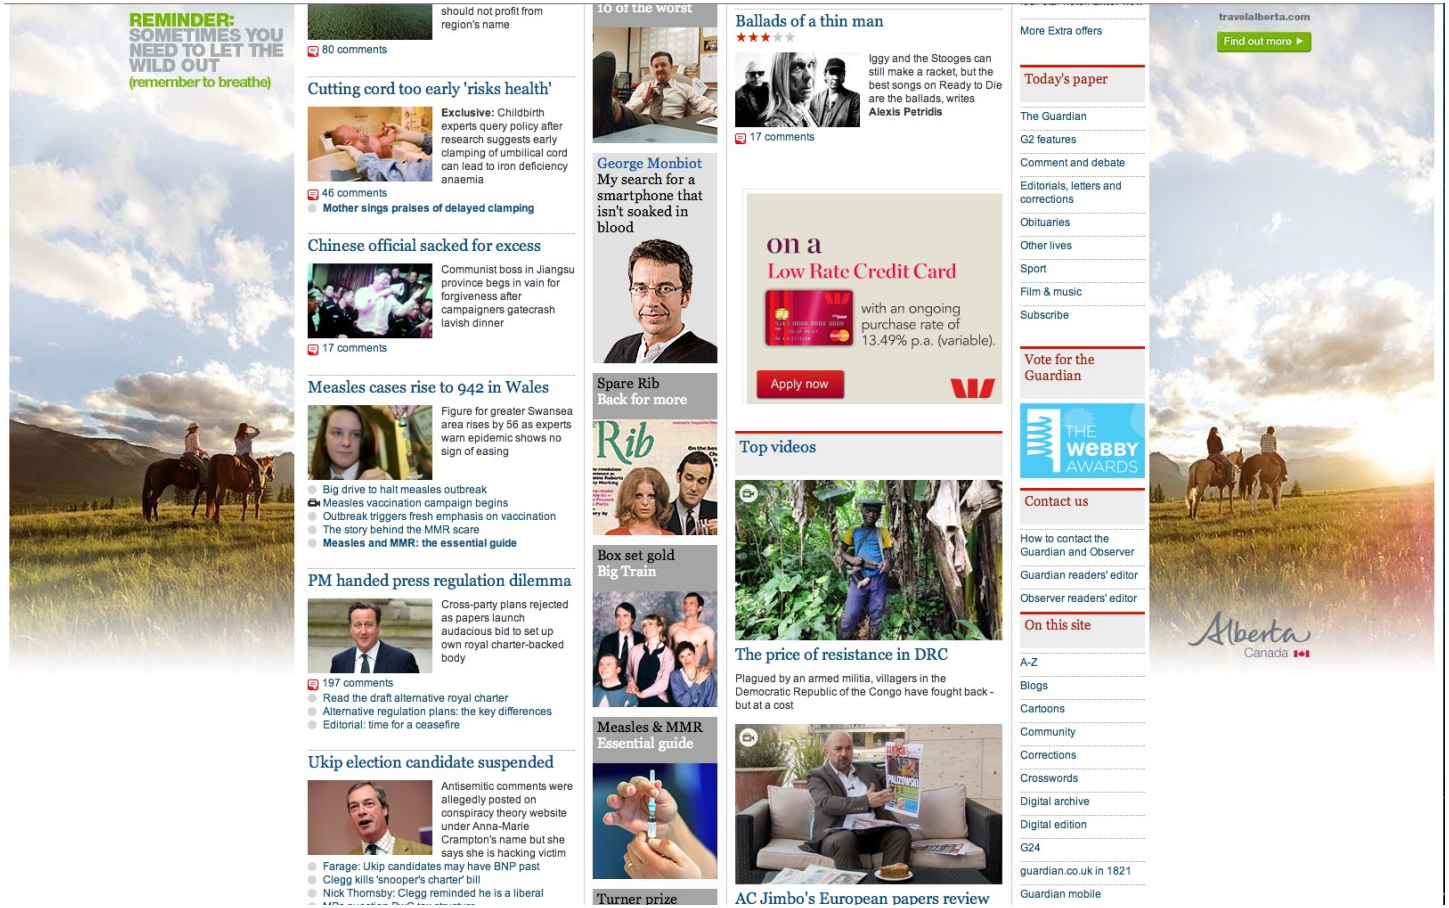
REMINDER:
SOMETIMES YOU NEED TO LET THE WILD OUT
(remember to breathe)
should not profit from region's name
80 comments
Cutting cord too early 'risks health'
Exclusive: Childbirth experts query policy after research suggests early clamping of umbilical cord can lead to iron deficiency anaemia
46 comments
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
George Monbiot
My search for a smartphone that isn't soaked in blood
Spare Rib
Back for more
Rib
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
travelalberta.com
Find out more
Alberta
Canada

Ads use Scripts

- Each time an ad is loaded the ad server loads creative content and scripts on to the client's browser
- The script can include action items to fetch 'network assets'
  - Typically used to load alternate images, sequences
  - Its not a generalized network stack, subject to constraints such as limited to certain object loads, reduced run-time library
- There are on-Load, on-Hover and on-Click actions
- We want to minimise interactions so we use on-Load scripting

This can work

- We can instrument the target host via an ad script
  - we can constrain the ad script to talk ONLY to our server(s)
  - And if we instrument the servers, we can infer the target host properties
- Ads try to deliver to new users all of the time
  - We want to measure new sample points all of the time to avoid implicit repeat bias in the measurement set
- Ads are biased towards ‘clicks’
  - We are not interested in clicks
  - We just want impressions
  - Impressions are far cheaper than clicks!

Advertising placement logic

- Fresh Eyeballs means a constant flow of unique endpoints
  - We have good evidence the advertising channel is able to sustain a constant supply of unique endpoints
- 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 endpoints 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

Advertising placement logic

- Set a ‘CPM’ bid in the ad
  - Clicks Per Millepressions: bid rate to pay per thousand impressions
- Uneven distribution of ads throughout the day
  - But we can compensate for this by running 24 x 1-hour campaigns per day
- Use multiple campaigns each with a constrained locale
  - That way we can ‘encourage’ the ad system to give the ad a broad placement
- Outcome: ~25M placements per day, on a mostly even placement model with end of day ‘soak’ to achieve budget goal

Advertising placement logic

- Se
- U
- U
- O
  m

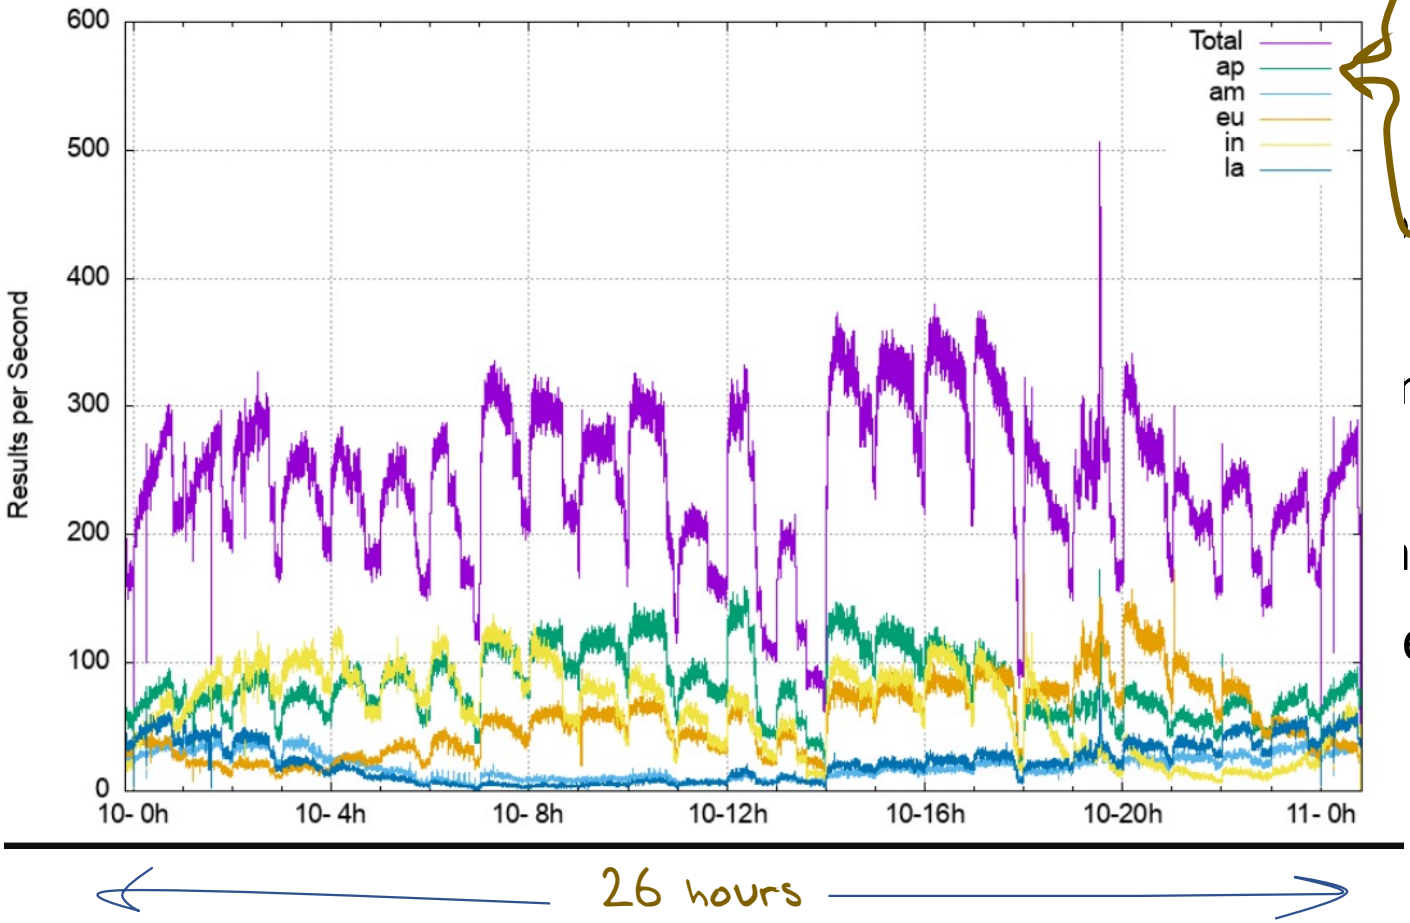

Asia (Singapore)
America (Dallas)
Europe (Frankfurt)
Asia (Mumbai)
S America (Sao Paulo)

ns per day

d placement

ement

What can be scripted in an Ad

Not much:

- http.FetchImg()

  i.e. attempt to retrieve a URL

But that's enough!

- It's EXACTLY what users do!
- A URL consists of a DNS question and an HTML question
- What if we point both the DNS and the HTML to use servers that we run?
- As long as each Ad execution uses unique DNS names, then we can push the user query back to our servers and avoid the use of caching
- We can't instrument the client, but we CAN (and do) instrument the server

Measuring IPv6 via Ads

- Use HTML5 code that is executed on ad impression
  - Client retrieves set of “tests” from an ad-controller
  - Client is given 10 URLs to load, including:
    - Dual Stack object
    - V4-only object
    - V6-only object
    - Result reporting URL (10 second timer)

    All DNS is dual stack

    All URLs use a unique DNS label

For Example:

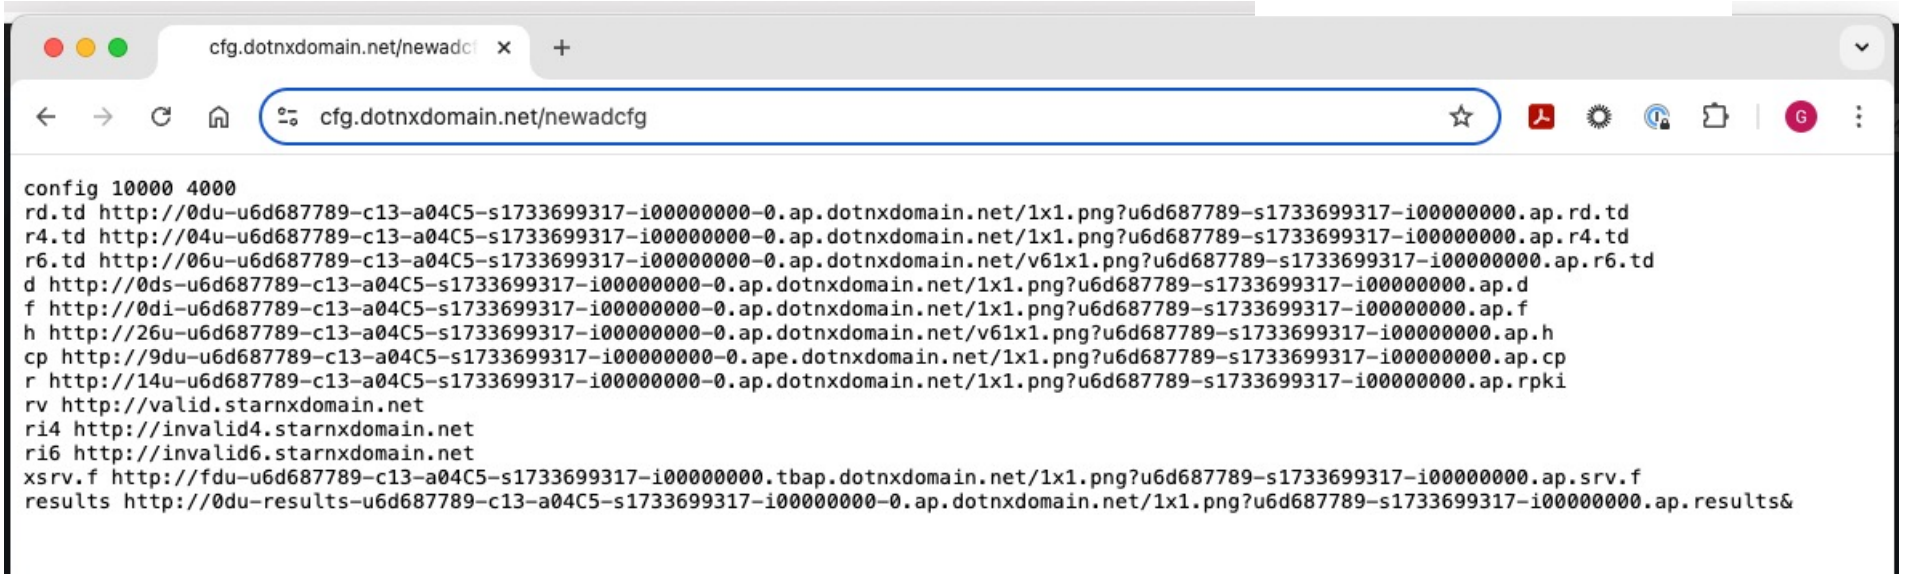
cfg.dotnxdomain.net/newadcf
cfg.dotnxdomain.net/newadcfg
config 10000 4000
rd.td http://0du-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.rd.td
r4.td http://04u-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.r4.td
r6.td http://06u-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/v61x1.png?u6d687789-s1733699317-i00000000.ap.r6.td
d http://0ds-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.d
f http://0di-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.f
h http://26u-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/v61x1.png?u6d687789-s1733699317-i00000000.ap.h
cp http://9du-u6d687789-c13-a04C5-s1733699317-i00000000-0.ape.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.cp
r http://14u-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.rpki
rv http://valid.starnxdomain.net
ri4 http://invalid4.starnxdomain.net
ri6 http://invalid6.starnxdomain.net
xsrv.f http://fdu-u6d687789-c13-a04C5-s1733699317-i00000000.tbap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.srv.f
results http://0du-results-u6d687789-c13-a04C5-s1733699317-i00000000-0.ap.dotnxdomain.net/1x1.png?u6d687789-s1733699317-i00000000.ap.results&

DNS Label Encoding

Think of a DNS name as a micro-coded instruction set directed to programmable DNS and HTTP servers …

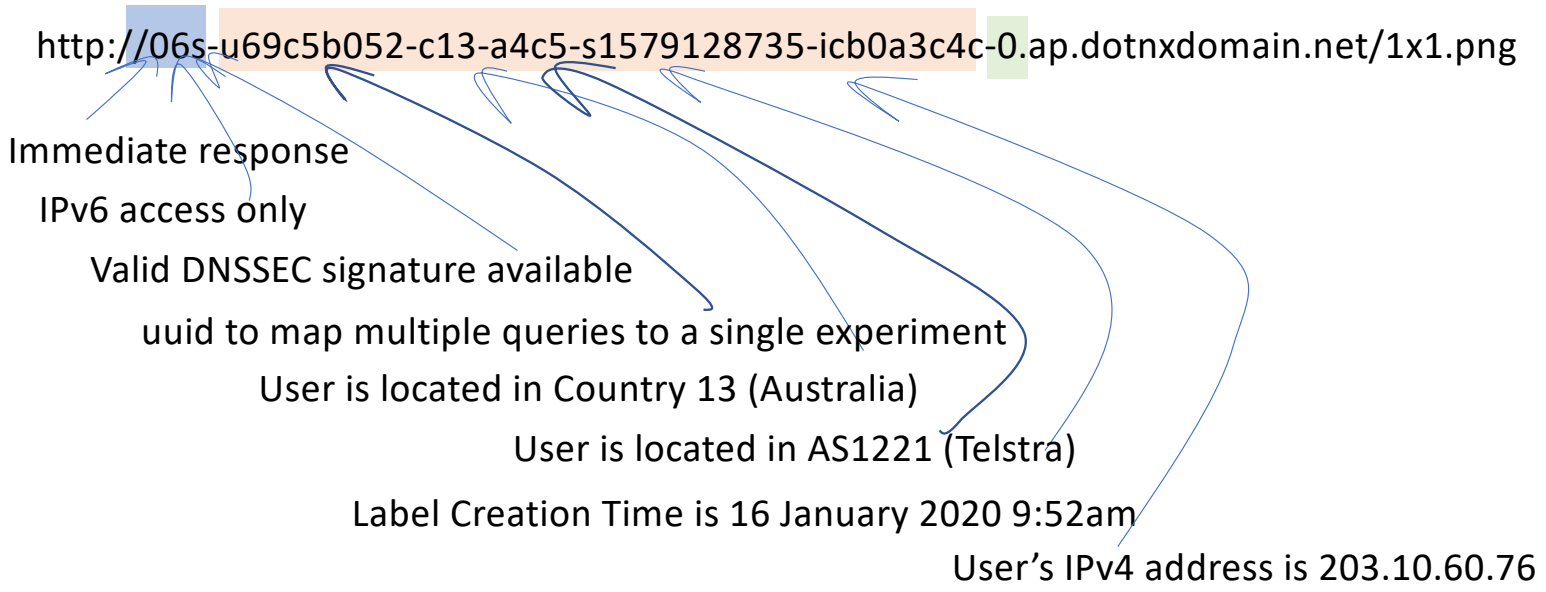

Experiment Server config

- There are six server sets, identically configured in VMs in DCs (Frankfurt, Singapore, Hong Kong, Dallas, Sao Paulo, Mumbai)
- The experiment script directs the client to the “closest” server set (based on geolocation of the client IP address)
- Server set has dedicated DNS and web content server VMs

Collected Data

Per Server, Per Day:

- HTTP access log
  (successfully completed fetches)
- DNS query log
  (incoming DNS queries)
- Packet capture
  All packets!

Data Analysis

For example – IPv6 measurement

- IPv6 “capable” means that the client successfully fetched the URL target that is only accessible using IPv6
- IPv6 “preferred” means that the client used IPv6 to fetch the dual stack URL target
- Aggregate data by origin AS and by geolocation CC
- “Normalise” the country data against estimates of national user populations (to compensate for aD placement bias at a national level)
- Generate IPv6 daily report and data to data set

V6 Time Series

Use of IPv6 for World (XA)

IPv6 Report

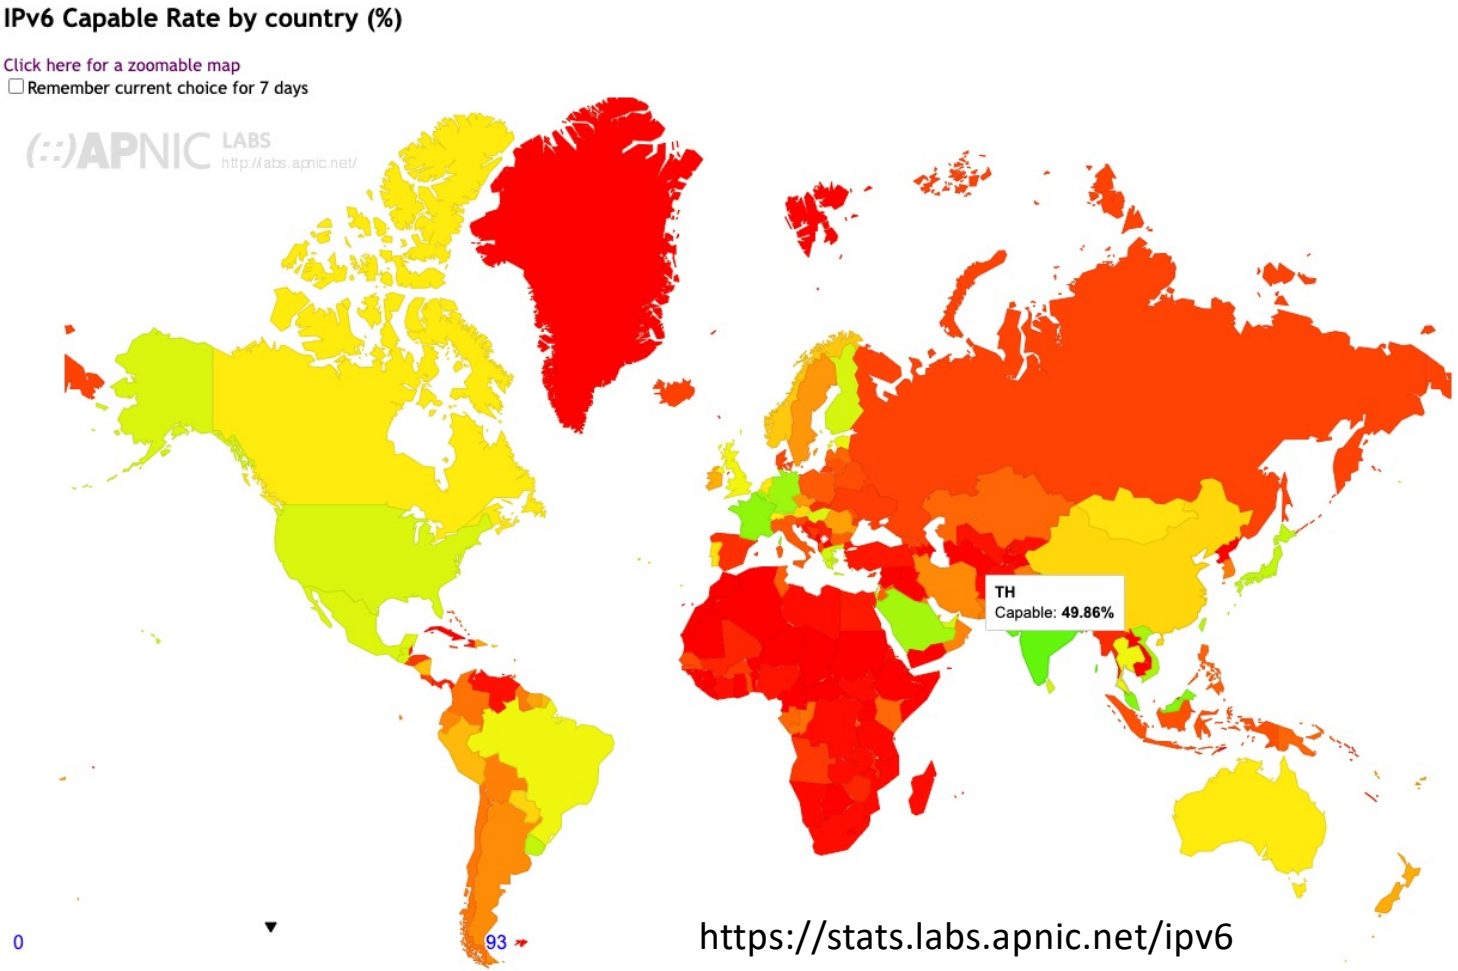

https://stats.labs.apnic.net/ipv6

IPv6 Report

IPv6 Capable Rate by country (%)

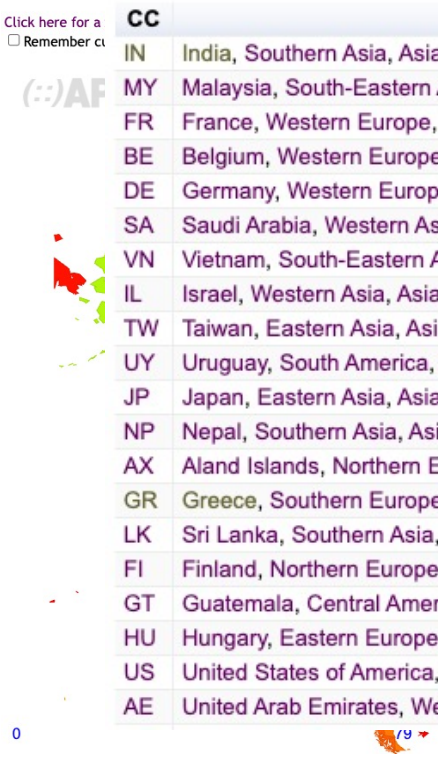

CC	Country	IPv6 Capable	IPv6 Preferred	Samples
IN	India, Southern Asia, Asia	79.81%	78.47%	148,742,179
MY	Malaysia, South-Eastern Asia, Asia	72.08%	68.38%	8,162,536
FR	France, Western Europe, Europe	70.39%	69.90%	16,054,034
BE	Belgium, Western Europe, Europe	68.33%	67.58%	2,541,184
DE	Germany, Western Europe, Europe	66.98%	66.28%	13,860,132
SA	Saudi Arabia, Western Asia, Asia	65.71%	64.40%	7,641,062
VN	Vietnam, South-Eastern Asia, Asia	64.23%	62.94%	9,891,922
IL	Israel, Western Asia, Asia	61.41%	54.37%	3,859,507
TW	Taiwan, Eastern Asia, Asia	60.94%	52.11%	4,127,683
UY	Uruguay, South America, Americas	59.32%	58.83%	621,842
JP	Japan, Eastern Asia, Asia	59.21%	55.38%	21,675,976
NP	Nepal, Southern Asia, Asia	58.09%	57.54%	2,101,637
AX	Aland Islands, Northern Europe, Europe	58.02%	57.56%	8,932
GR	Greece, Southern Europe, Europe	56.75%	56.46%	2,289,524
LK	Sri Lanka, Southern Asia, Asia	56.46%	55.76%	1,298,790
FI	Finland, Northern Europe, Europe	55.23%	54.54%	1,524,989
GT	Guatemala, Central America, Americas	55.10%	54.33%	628,593
HU	Hungary, Eastern Europe, Europe	54.57%	53.98%	2,638,417
US	United States of America, Northern America, Americas	54.33%	53.29%	92,238,748
AE	United Arab Emirates, Western Asia, Asia	54.16%	53.16%	1,598,869

https://stats.labs.apnic.net/ipv6

Per-Country Time Series - India

https://stats.labs.apnic.net/ipv6/IN

Per-Country Time Series - India

ASN	AS Name	IPv6 Capable	IPv6 Preferred	Samples
AS55836	RELIANCEJIO-IN Reliance Jio Infocomm Limited	95.31%	93.74%	70,668,475
AS45609	BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service	91.25%	89.97%	36,257,071
AS24560	AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services	65.31%	63.36%	8,388,823
AS38266	VIL-AS-AP Vodafone Idea Ltd	89.34%	88.34%	6,559,337
AS45271	ICLNET-AS-AP Idea Cellular Limited	87.88%	86.19%	4,469,821
AS9829	BSNL-NIB National Internet Backbone	4.24%	4.03%	3,193,236
AS24309	CABLELITE-AS-AP Atria Convergence Technologies Pvt. Ltd. Broadband Internet Service Provider INDIA	38.04%	37.67%	1,077,452
AS133982	EXCITEL-AS-IN Excitel Broadband Private Limited	0.38%	0.09%	968,656
AS45916	GTPL-AS-AP Gujarat Telelink Pvt Ltd	37.21%	34.54%	858,588
AS138754	KVBPL-AS-IN Kerala Vision Broad Band Private Limited	17.93%	17.67%	792,974
AS133661	NETPLUS-AS Netplus Broadband Services Private Limited	50.35%	47.17%	782,841
AS17488	HATHWAY-NET-AP Hathway IP Over Cable Internet	0.43%	0.22%	719,382
AS24186	RAILTEL-AS-IN RailTel Corporation of India Ltd	8.18%	7.76%	548,135
AS17665	ONEBROADBAND ONEOTT INTERTAINMENT LIMITED	3.30%	2.95%	541,747
AS55577	CABLELITE-AS-AP Atria Convergence Technologies Ltd.	26.72%	26.35%	495,761

https://stats.labs.apnic.net/ipv6/IN

Per-Network Time Series

IPv6 Per-Country Deployment for AS55836: RELIANCEJIO-IN Reliance Jio Infocomm Limited, India (IN)

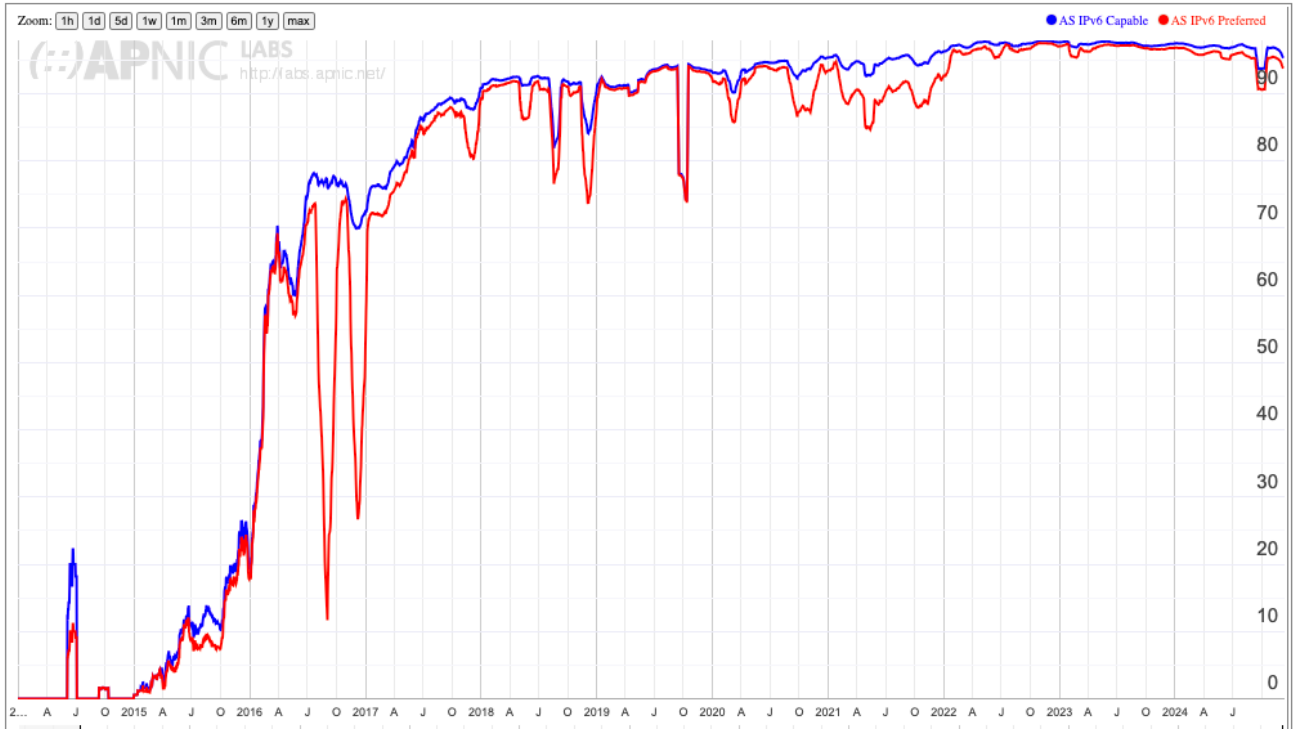

https://stats.labs.apnic.net/ipv6/AS55836

What about DNSSEC Use?

Can we use the same platform to measure the proportion of users who sit behind DNS resolvers that perform DNSSEC validation?

What about DNSSEC Use?

- Can we use the same platform to measure the proportion of users who sit behind DNS resolvers that perform DNSSEC validation?

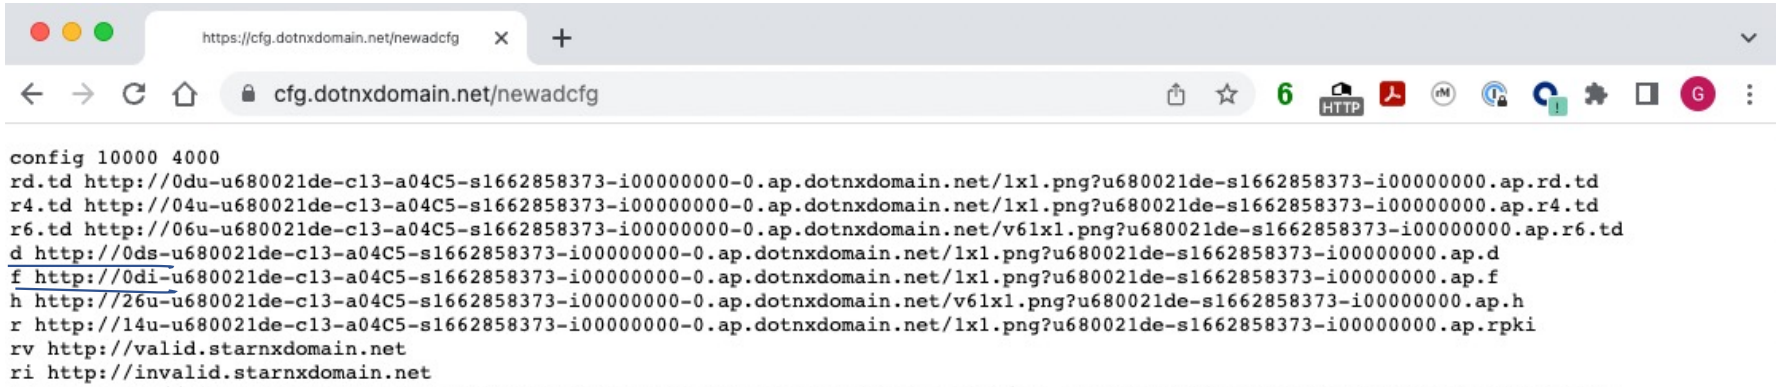

DNSSEC Test

- A client is sitting behind DNSSEC validating resolvers if:
  - It fetches the validly-signed DNS URL
  - It does not fetch the invalidly-signed DNS URL
  - It queries for the address using EDNS0 DNSSEC OK field
  - It queries for the zone DNSKEY record for both DNS names *
- A client "partially" validates if the validation failure causes the client to use an alternative non-validating resolver and resolve the name

* This last one is a bit of a challenge as we want to see these queries at the authoritative server and not have them masked by caching – so we use a LDNS-based auth server to create a dynamic DNS server that constructs a DNSSEC signed delegated zone on the fly

DNSSEC Results

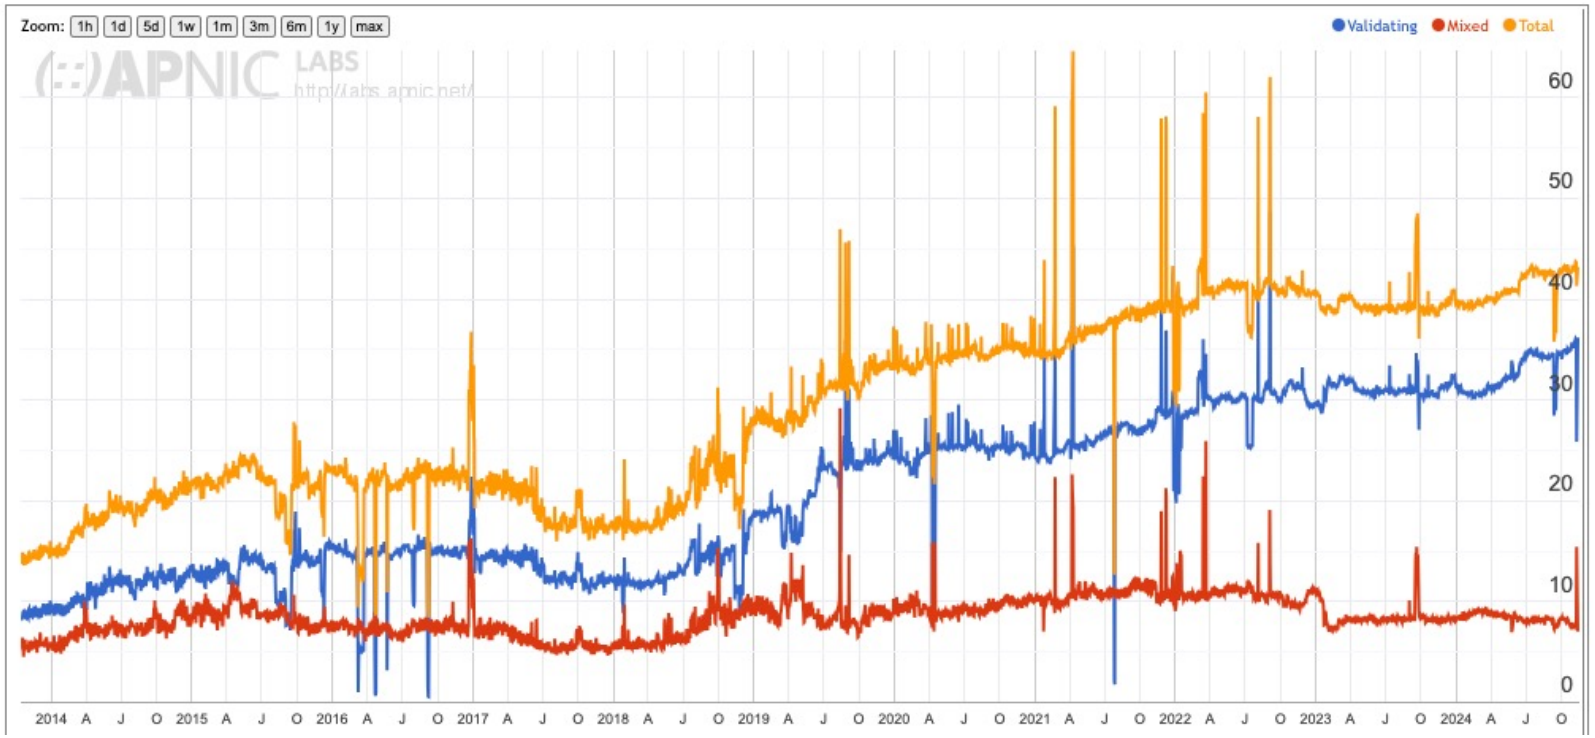

https://stats.labs.apnic.net/dnssec/XA

DNSSEC Results

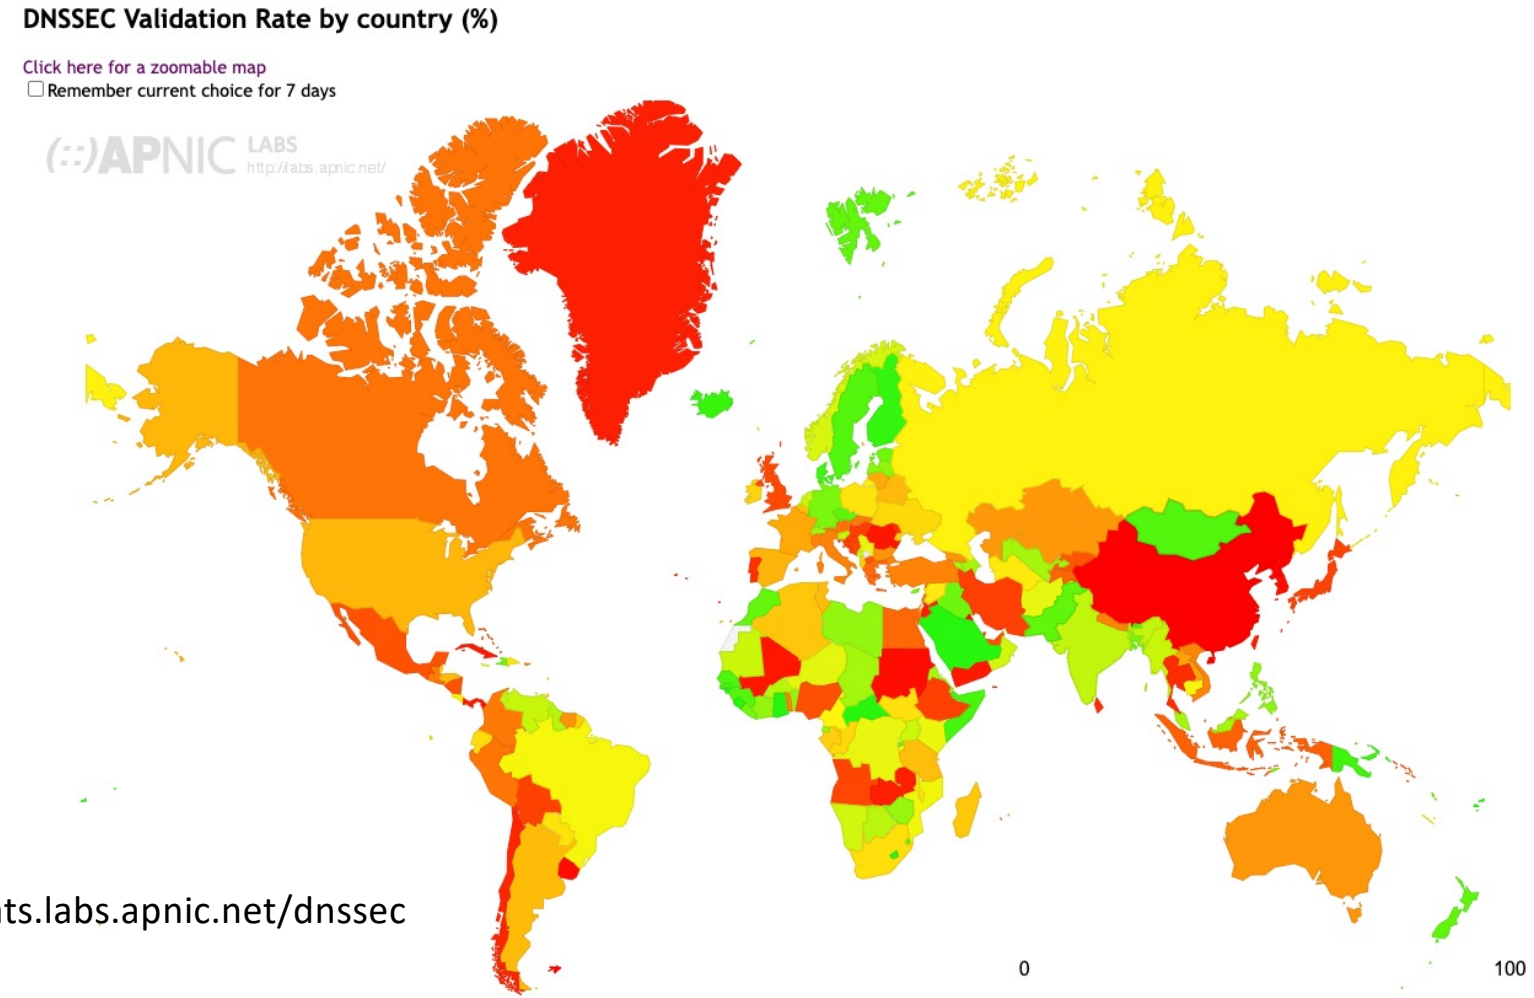

https://stats.labs.apnic.net/dnssec

DNSSEC Validation in India

ASN	AS Name	DNSSEC Validates	Partial Validation	Samples
AS55836	RELIANCEJIO-IN Reliance Jio Infocomm Limited	99.75%	0.19%	1,988,476
AS45609	BHARTI-MOBILITY-AS-AP Bharti Airtel Ltd. AS for GPRS Service	5.32%	2.27%	955,922
AS24560	AIRTELBROADBAND-AS-AP Bharti Airtel Ltd., Telemedia Services	2.68%	1.69%	204,627
AS38266	VIL-AS-AP Vodafone Idea Ltd	80.48%	2.43%	164,999
AS45271	ICLNET-AS-AP Idea Cellular Limited	59.47%	3.93%	119,917
AS9829	BSNL-NIB National Internet Backbone	13.51%	1.92%	69,455
AS24309	CABLELITE-AS-AP Atria Convergence Technologies Pvt. Ltd. Broadband Internet Service Provider INDIA	94.37%	3.43%	26,785
AS133982	EXCITEL-AS-IN Excitel Broadband Private Limited	89.86%	10.00%	23,671
AS45916	GTPL-AS-AP Gujarat Telelink Pvt Ltd	22.30%	41.44%	22,358
AS133661	NETPLUS-AS Netplus Broadband Services Private Limited	17.56%	80.72%	19,356
AS17488	HATHWAY-NET-AP Hathway IP Over Cable Internet	26.39%	27.86%	16,139
AS24186	RAILTEL-AS-IN RailTel Corporation of India Ltd	61.97%	34.06%	15,657
AS138754	KVBPL-AS-IN Kerala Vision Broad Band Private Limited	99.05%	0.80%	14,535

https://stats.labs.apnic.net/dnssec/IN

Other server side measurement techniques

We treat the DNS name (and the full URL) as a set of server instructions and use a combination of dynamic DNS, NGINX modules and 2-step packet processing to generate specific server-side behaviours that we want to measure

Other server-side measurement techniques - Glueless DNS

"Glueless" DNS delegation to provide explicit confirmation that a resolver has received a response

- The client is forced to resolve the name of the zone name servers before proceeding with the original resolution task
- If the name of the name servers is dynamically generated and unique then DNS caching won't help
- We've used this technique to measure:
  - IDN support
  - qname minimisation
  - DNS fragmentation management
  - DNS dual stack behaviour
  - Recursive Resolvers' server selection

Other server side measurement techniques - Repeats

Explicitly directed repeat fetches to trigger content directives for HTTP/3

- Chrome relies on receiving an Alt-Svc: content directive before it will switch over to use HTTP for fetches from this server
- Which means it will only use HTTP/3 on the second (or subsequent) fetch
- To create this behaviour we have to allow the measurement script to schedule this fetch multiple times, with a small idle interval between fetches

Other server side measurement techniques - SERVFAIL DNS

SERVFAIL responses to force the client resolver to cycle through all configured recursive resolvers

- This technique uses a LDNS path that responds with a SERVFAIL response for all DNS queries for this name set
- This is intended to cause the local resolver to cycle through all locally configured recursive resolvers to find a resolver that will respond to the query
- We used this is exposing the set of resolvers that a user may use to resolve a name

Other server side measurement techniques - IPv6 Packet Mangling

- IPv6 packet manipulation to insert crafted IPv6 packets into an established TCP stream
  - To manipulate an IPv6 packet to manually control fragmentation and other forms of Extension Headers we use a “front end” unit to pick up incoming packets and pass them to a conventional back end server
  - The return packet is modified to add the appropriate Extension Header and/or Fragment before passing back to the end client

Measurement Projects

- IPv6 Performance (connection reliability and relative speed)
- IPv6 Fragmentation
- IPv6 Extension Header loss Rates (HBH and DST)
- DNS: Use of ECDSA and EDDI DNSSEC signing algorithms
- DNS: Fragmentation Drop (and TCP support)
- DNS resolver use profile (use of open DNS resolvers)
- DNS KSK roll probes (RFC8509)
- Support for QUIC use (HTTP/3)
- Support for Route Origination Validation
- Zombies and tracking

https://stats.labs.apnic.net

Server Side Measurement

- This approach complements client side measurements (CAIDA's ARK, RIPE NCC's Atlas) and network-level internal measurements by using a large scale server side measurement platform
- In this form of server-side measurement the client does what clients always do - fetch URLs
- We can test particular client behaviours and network behaviours by deliberately altering the server-side behaviour and triggering the behaviour in a measured behaviour
- The benefit of this approach is that rather than measuring the effect and inferring the cause, in this approach we trigger a cause and then correlate the observed outcomes against the known cause.

Thanks!

Measurement Reports at APNIC Labs:
https://stats.labs.apnic.net