



# Measuring the End User

Geoff Huston  
APNIC Labs

# Measurement Bias

When we first looked at measuring in the Internet, it was all about the network, and the distinction between network management and network measurement was not very clear

We ended up measuring what's easy to measure and often missed measuring what's useful to understand

# "Measurable" Questions?

- How many routes are IPv6 routes?
- How many service providers offer IPv6?
- How many domain names have AAAA RRs?
- How many domains are DNSSEC signed?
- How many DNS queries are made over IPv6?
- How much traffic uses IPv6?
- How many connections use IPv6?

...

# Users vs Infrastructure

None of these specific measurement questions really embrace the larger questions about the end user experience

They are all aimed at measuring an aspect of of behaviour within particular parameters of the network infrastructure, but they don't encompass how the end user assembles a coherent view of the network





The Internet is all about  
US!



# What's the question?

*How many **users** do <x>?*

- How many **users** can are running IPv6?
- How many **users** are using DNSSEC validation?
- How many **users** support ECDSA in digital signatures in DNSSEC?
- How many **users** can resolve a DNS name?  
etc

# Private vs Public

- Very few measurements on the Internet are public
- Most “all of Internet” metrics are wild-eyed guesses
  - How many people use the Internet?
  - How many devices use the Internet
  - How much traffic is passed across the Internet?
- And the bits that aren't guesses are often folded into proprietary data

# The Challenge:

How can we undertake meaningful public measurements that quantify aspects of users' experiences drawn from across the entire Internet that does not rely on access to private data?

# For example... IPv6

- It would be good to know how we are going with the transition to IPv6
- And it would be good everyone to know how everyone else is going with the transition to IPv6
- What **can** we measure?
  - IPv6 in the DNS – AAAA records in the Alexa top N
  - IPv6 in routing – IPv6 routing table
  - IPv6 traffic exchanges – traffic graphs
- What **should** we measure?
  - How many connected devices on today's Internet are capable of making IPv6 connections?

How to measure millions of  
end devices for their IPv6  
capability?



# How to measure millions of end devices for their IPv6 capability?

a) Be



How to measure millions of  
end devices for their IPv6  
capability?

a) Be Google

OR

b) Have your measurement code run on a  
million end devices

# 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

travelberta.com

Find out more

## 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




# Ads are ubiquitous

The image shows a screenshot of a news website with a large red hand-drawn circle highlighting several areas. The highlighted areas include:

- A large advertisement on the left side of the page featuring a landscape with people on horseback and the text: "REMINDER: SOMETIMES YOU NEED TO LET THE WILD OUT (remember to breathe)".
- A "Low Rate Credit Card" advertisement in the center, featuring a red credit card and the text: "On a Low Rate Credit Card with an ongoing purchase rate of 13.49% p.a. (variable). Apply now".
- A "Ballads of a thimble" advertisement at the top right, featuring a photo of a man and the text: "Iggy and the Stooges can still make a racket, but the best songs on Ready to Die are the ballads, writes Alexis Petridis".
- A "Spare Rib Back for more" advertisement featuring a photo of a woman and the text: "Box set gold Big Train".
- A "Measles & MMR Essential guide" advertisement featuring a photo of a hand holding a syringe.
- A "Turner prize" advertisement featuring a photo of a man sitting on a sofa.
- A "AC Jimbo's European papers review" advertisement featuring a photo of a man sitting on a sofa.

The background shows various news articles and a sidebar with navigation links. The sidebar includes links such as "More Extra offers", "Today's paper", "The Guardian", "G2 features", "Comment and debate", "Editorials, letters and corrections", "Opinion", "Other news", "Sport", "Film & music", "Subscriptions", "Vote for the Guardian", "THE WOBBY 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", and "Guardian mobile".

# Ads are ubiquitous



# Ads use active scripts

- Advertising channels use active scripting to make ads interactive
  - This is not just an ‘animated gif’ – it uses a script to sense mouse hover to change the displayed image



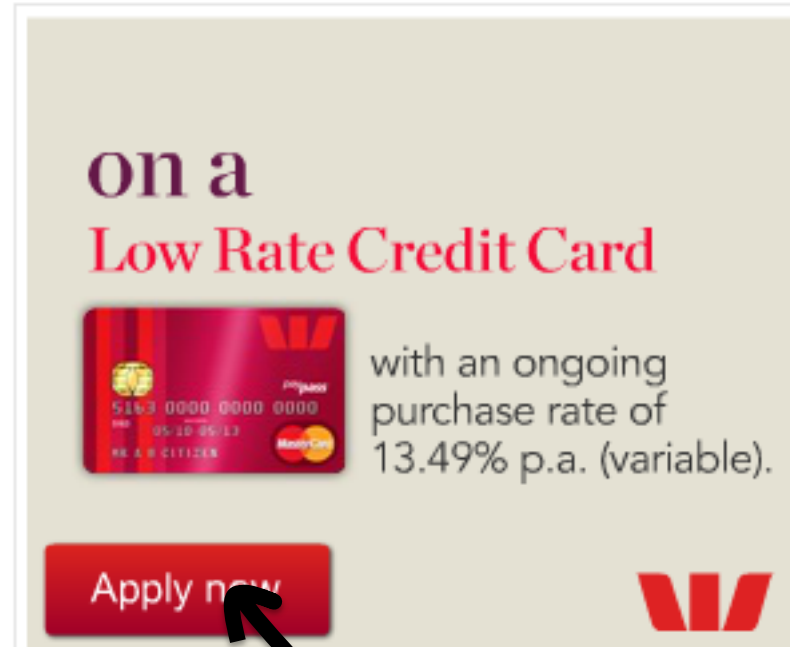
0%  
p.a.  
on purchases

\*New cards only.  
Conditions apply.


Apply now



A black arrow points to the 'Apply now' button.




on a  
Low Rate Credit Card



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

Apply now



A black arrow points to the 'Apply now' button.



# Adobe 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 over what connections can be made
- Flash has asynchronous 'threads' model for event driven, sprite animation

# Adobe Flash and the network

- Flash includes primitives in 'actionscript' to fetch 'network assets'
  - Flash is used to load alternate images, sequences
  - So these days we use HTML5 as the vehicle for the measurement script
- Flash has asynchronous threads for event driven, sprite animation

# APNIC's measurement technique

- Craft a script which fetches URLs to measure.
  - URLs are reduced to a notional '1x1' image which is not added to the browser's display manager and is not displayed
  - URLs trigger DNS resolution via whatever name resolution mechanism is used by the local browser and host
  - We encode data transfer from the client to the server in the name of fetched URLs
    - Could use the DNS as the information conduit:
      - Result is returned by DNS name
    - Could use HTTP as the information conduit
      - Result is returned via parameters attached to an HTTP GET command
- We use a combination of http requests and server logs

# The Ad Measurement Technique



Ad Server



End user

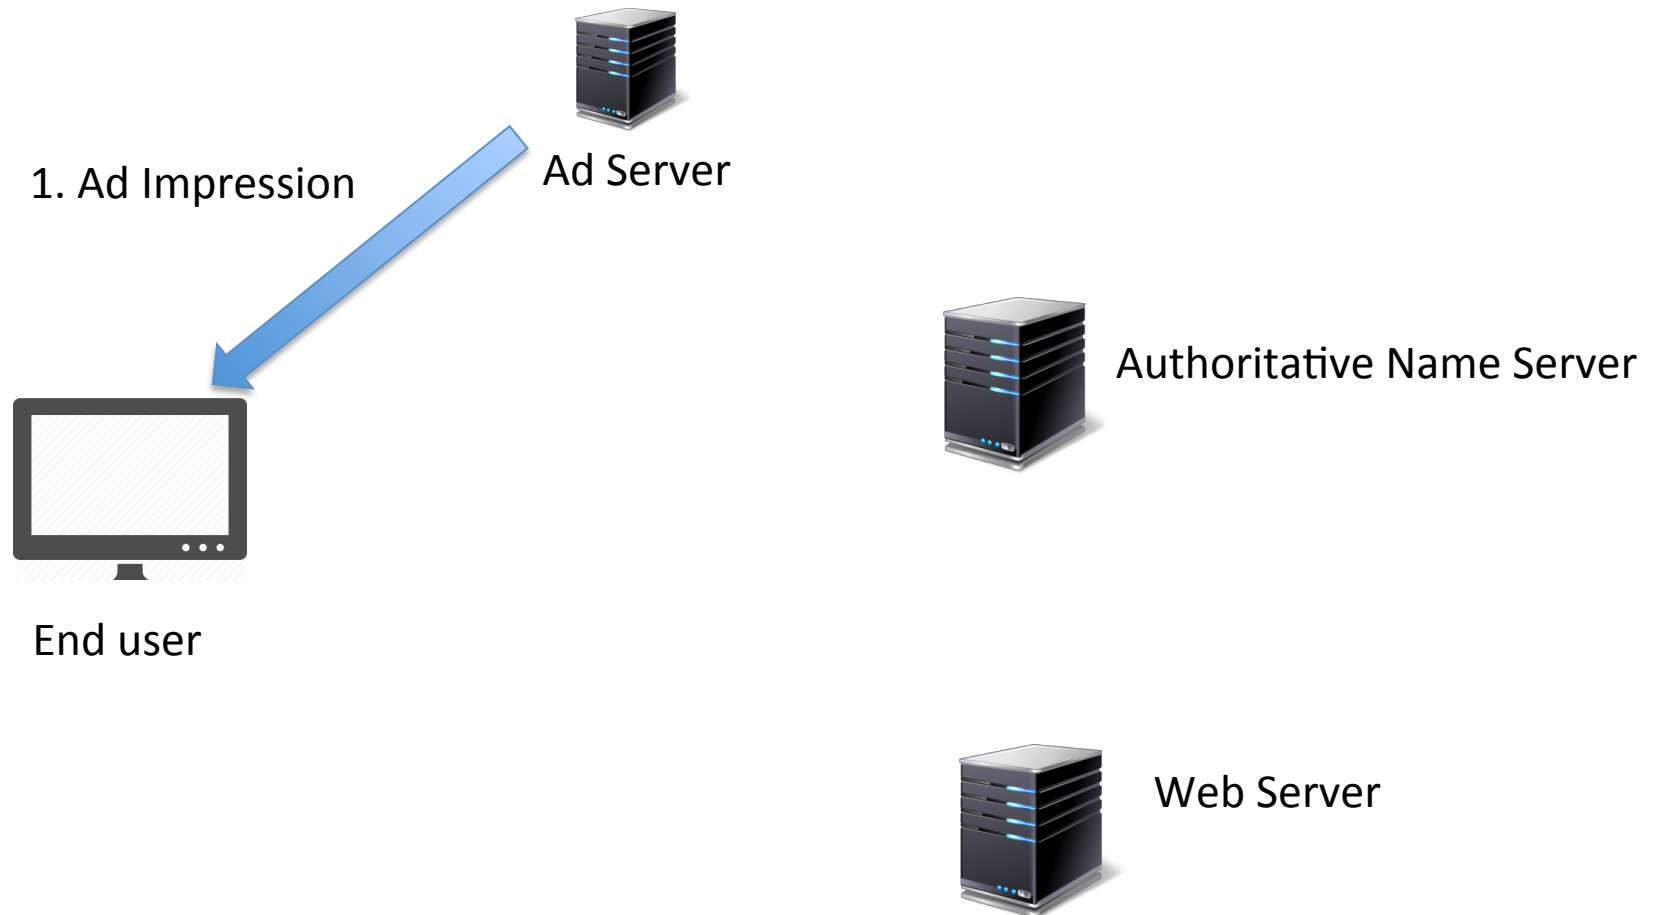


Authoritative Name Server

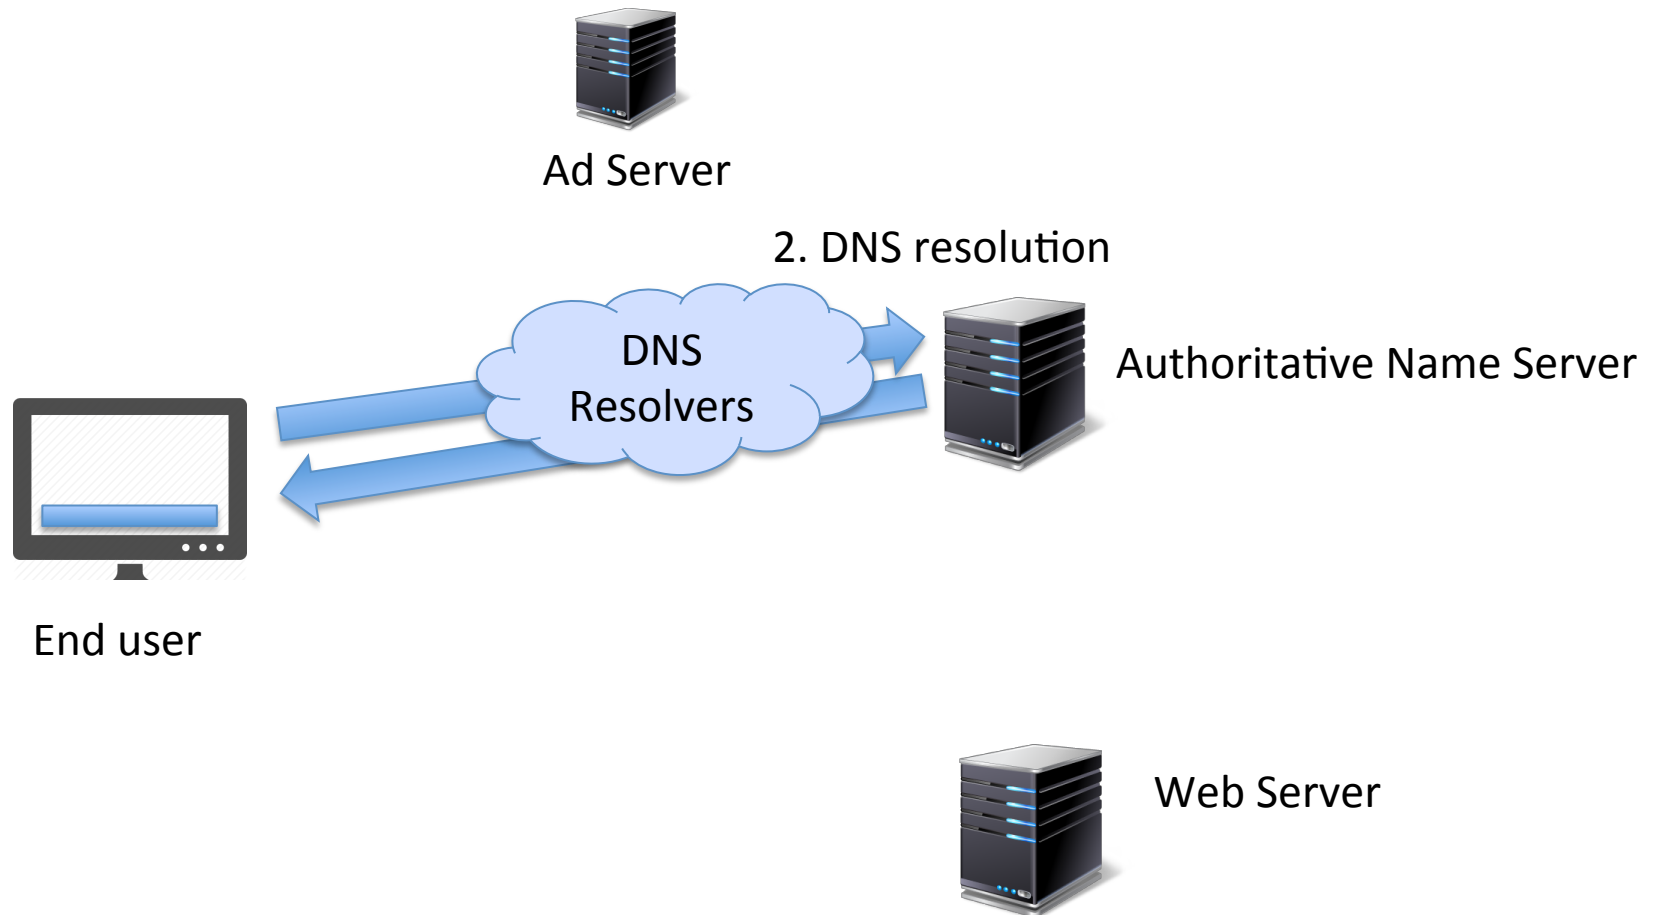


Web Server

# The Ad Measurement Technique

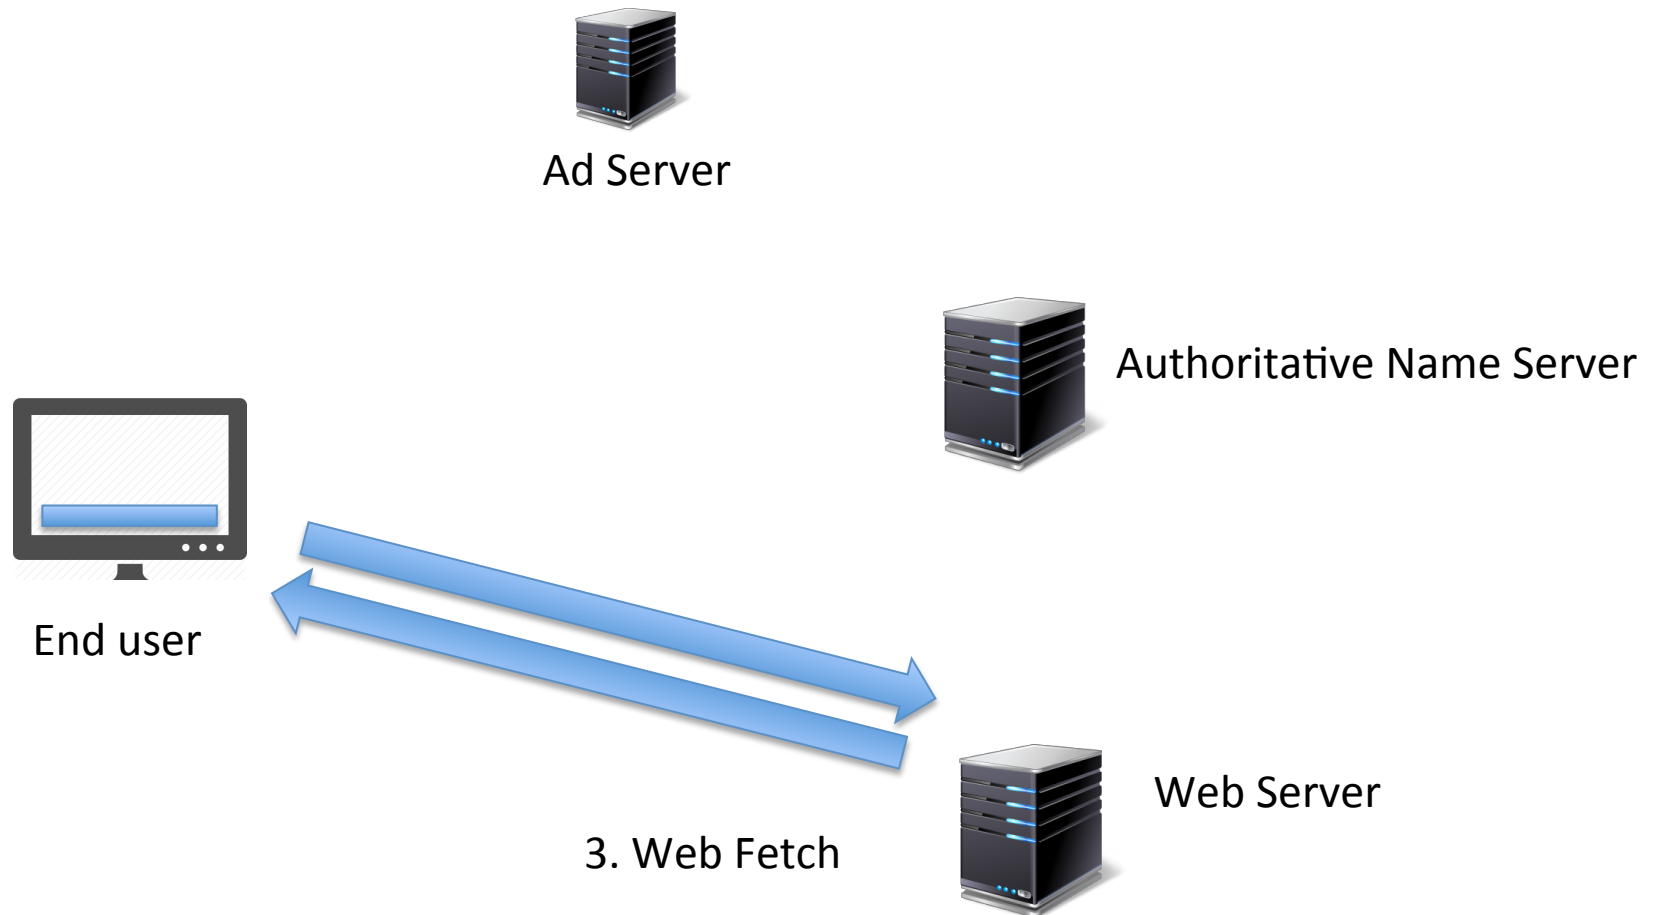


# The Ad Measurement Technique

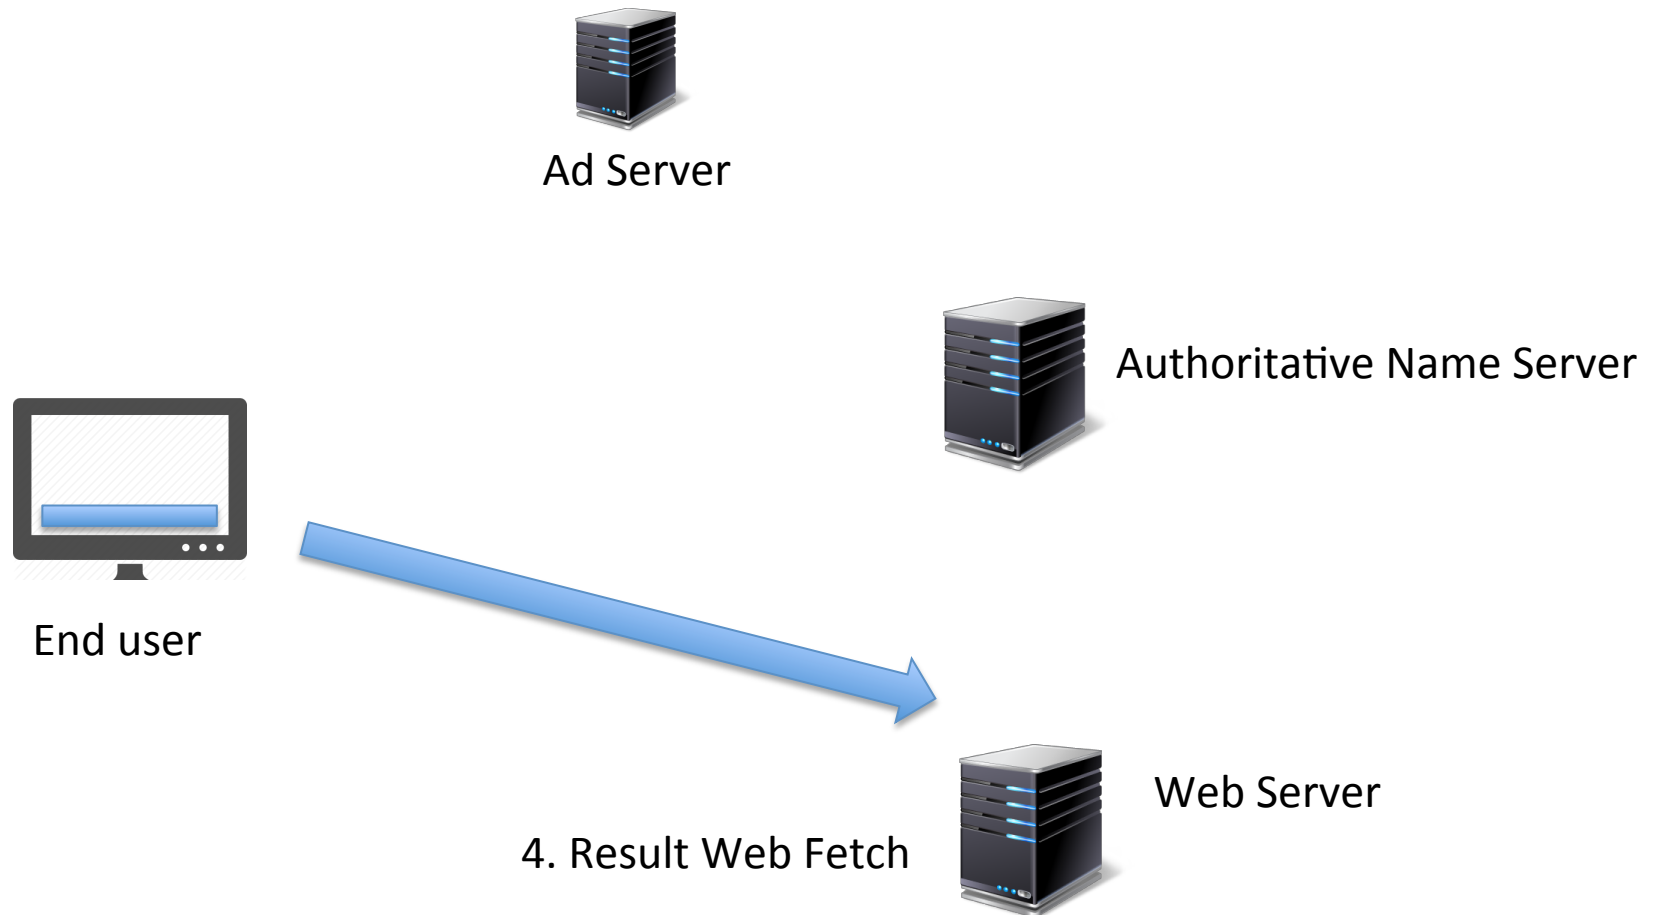




# The Ad Measurement Technique




# The Ad Measurement Technique




# 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 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 the ad placement machinery an algorithm tries hard to achieve even placement loads, but in the end, will 'soak' place your ad to achieve enough views to bill you that target of \$100


# Ad Placement Training - Day 1



# Ad Placement Training - Day 2




# Ad Placement Training - Day 3






# Ad Placement Training - Day 4



# Ad Placement Training - Days 5, 6 & 7



# Fresh Eyeballs



# Success!

- 2.5M – 3M samples per day – mostly new!
- Large sample space across much of the known Internet
- Assemble a rich data set of end user addresses and DNS resolvers

# Success ... of a sort!

- What we are after is a random sample of the entire Internet
- And we are close
- But what we have is a data set biased towards “cheap” eyeballs in fixed networks

# "Raw" AD counts per day

155,430	VN Vietnam
103,517	CN China
92,107	MX Mexico
79,092	TH Thailand
73,702	IN India
65,402	PK Pakistan
64,121	BR Brazil
54,637	TR Turkey
52,532	US United States of America
52,240	AR Argentina
48,315	CO Colombia
45,216	ID Indonesia
39,839	PE Peru
36,962	RU Russian Federation
34,529	PH Philippines
33,899	EG Egypt
22,983	TW Taiwan
22,712	RO Romania
22,490	UA Ukraine
22,403	ES Spain

IP address to country code mapping for  
experiments placed on the 24<sup>th</sup> May 2015

# ITU-T's Internet User Census

155,430	VN Vietnam	668,493,485	China
103,517	CN China	282,384,872	United States of America
92,107	MX Mexico	252,482,905	India
79,092	TH Thailand	110,345,878	Brazil
73,702	IN India	109,390,190	Japan
65,402	PK Pakistan	87,305,661	Russian Federation
64,121	BR Brazil	72,663,301	Nigeria
54,637	TR Turkey	71,823,404	Indonesia
52,532	US United States of America	71,174,958	Germany
52,240	AR Argentina	61,579,582	Mexico
48,315	CO Colombia	57,306,333	United Kingdom of Great Britain and Northern Ireland
45,216	ID Indonesia	54,114,094	France
39,839	PE Peru	45,416,941	Iran (Islamic Republic of)
36,962	RU Russian Federation	45,019,465	Egypt
34,529	PH Philippines	42,187,842	Republic of Korea
33,899	EG Egypt	41,780,667	Philippines
22,983	TW Taiwan	40,980,368	Vietnam
22,712	RO Romania	39,256,999	Bangladesh
22,490	UA Ukraine	35,793,673	Italy
22,403	ES Spain	35,503,461	Turkey

ITU's estimates of number of internet users per country


# "Weighting" sample data to correct AD Placement bias

We "weight" the raw data by:

- Geolocating the IP address to a particular country
- Multiplying the sample by the relative weight of the country



# Weighting the Results



CC	Country	DNSSEC Validates	Uses Google PDNS	Samples	Weight	Weighted Samples
CN	China, Eastern Asia, Asia	3.75%	8.84%	6726432	2.62	17603898
US	United States of America, Northern America, Americas	22.71%	10.05%	2654162	2.8	7436237
IN	India, Southern Asia, Asia	9.60%	21.22%	2788239	2.38	6648806
BR	Brazil, South America, Americas	25.96%	19.77%	3654367	0.8	2905814
JP	Japan, Eastern Asia, Asia	7.31%	4.77%	421526	6.83	2880651
RU	Russian Federation, Eastern Europe, Europe	14.19%	10.36%	2331208	0.99	2299084
NG	Nigeria, Western Africa, Africa	13.75%	42.50%	33555	57.03	1913494
ID	Indonesia, South-Eastern Asia, Asia	13.69%	14.46%	2399551	0.79	1891377
DE	Germany, Western Europe, Europe	16.66%	4.73%	437365	4.29	1874295
MX	Mexico, Central America, Americas	4.78%	8.82%	3852726	0.42	1621619
GB	United Kingdom of Great Britain and Northern Ireland, Northern Europe, Europe	7.25%	6.33%	609639	2.48	1509090
FR	France, Western Europe, Europe	24.23%	3.59%	1067161	1.34	1425021
IR	Iran (Islamic Republic of), Southern Asia, Asia	21.05%	37.95%	5895	202.88	1195992
EG	Egypt, Northern Africa, Africa	14.86%	16.53%	1478598	0.8	1185526
KR	Republic of Korea, Eastern Asia, Asia	1.78%	2.56%	871624	1.27	1110963
PH	Philippines, South-Eastern Asia, Asia	10.93%	12.62%	1360250	0.81	1100239
VN	Vietnam, South-Eastern Asia, Asia	29.67%	46.79%	5580740	0.19	1079161
BD	Bangladesh, Southern Asia, Asia	31.78%	42.05%	459167	2.25	1033783
IT	Italy, Southern Europe, Europe	15.78%	19.39%	805477	1.17	942582

# Measuring ALL of the Internet

It's not perfect by any means, but it is a reasonable first pass to correct for the implicit ad placement bias in the raw data

So now we have a method to measure a sample of Internet users and a process that can relate that measurement back to the Internet as a whole.

How can we use this?

# What does this allow?

In providing an end user with a set of URLs to retrieve we can examine:

- Protocol behaviour

  - e.g.: V4 vs V6, protocol performance, connection failure rate

- DNS behaviours

  - e.g.: DNSSEC use, DNS resolution performance, DNS response size, crypto protocol performance,...

# 1. Measuring IPv6

# Measuring IPv6

Client is given 4 unique URLs to load:

- Dual Stack object
- V4-only object
- V6-only object
- Result reporting URL (10 second timer)

We want to compare the number of end devices that can retrieve the V6-only object to the number of devices that can retrieve the V4-only object (V6 Capable)

We can also look at the number of end devices that use IPv6 to retrieve the Dual Stack Object (V6 Preferred)

# What we see (Web Log)

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:29 +0000] "GET /newadcfg/ad.py?A=2121&N&R&F HTTP/1.1" 200 799 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 u281fd425-s1438646489 1438646489.894 cfg.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.rd.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.e HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r6.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.578 O6u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.f HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.871 Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.d HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.159 Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r4.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.448 O4u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.results&zrtd-390.zr4td-1548.zr6td-678.zd-1258.ze-390.zf-971. HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.815 Odu-results-u281fd425-x-i5097.ap.dotnxdomain.net

# What we see (Web Log)

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:29 +0000] "GET /newadcfg/ad.py?A=2121&N&R&F HTTP/1.1" 200 799 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 u281fd425-s1438646489 1438646489.894 cfg.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.rd.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

This is a Mac OSX system, using OS X 10.9.5, with Chrome 44.0.2403.125

ap.e HTTP/1.1" 200 68  
eWebKit/537.36 (KHTML,

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r6.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.578 O6u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.f HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.871 Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.d HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.159 Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r4.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.448 O4u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.results&zrtd-390.zr4td-1548.zr6td-678.zd-1258.ze-390.zf-971. HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.815 Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net



# What we see (Web Log)

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:29 +0000] "GET /newadcfg/ad.py?A=2121&N&R&F HTTP/1.1" 200 799 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 u281fd425-s1438646489 1438646489.894 cfg.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.rd HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.e HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r6 HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.578 O6u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.f HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.871 Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temo "http like G This system can do IPv6, and prefers to use IPv6 in dual stack contexts TP/1.1" 200 68 it/537.36 (KHTML,

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r4.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.448 O4u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.results&zrtd-390.zr4td-1548.zr6td-678.zd-1258.ze-390.zf-971. HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.815 Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

# What we see (Web Log)

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:29 +0000] "GET /newadcfg/ad.py?A=2121&N&R&F HTTP/1.1" 200 799 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 u281fd425-s1438646489 1438646489.894 cfg.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.rd.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

This experiment ran through to completion

temora.ra "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r6.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.578 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.f HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.871 Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.d HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.159 Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r4.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.448 O4u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.results&zrtd=390.zr4td-1548.zr6td-678.zd-1258.ze-390.zf-971. HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.815 Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

# What we see (Web Log)

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:29 +0000] "GET /newadcfg/ad.py?A=2121&N&R&F HTTP/1.1" 200 799 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 u281fd425-s1438646489 1438646489.894 cfg.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.rd.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.e HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.290 Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

This user is a customer of TMNET in Malaysia, AS4788

HTTP/1.1" 200 68  
537.36 (KHTML,

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:30 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.f HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646490.871 Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net


temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.d HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.159 Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 124.13.125.185 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.r4.td HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.448 O4u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net

temora.rand.apnic.net 2001:e68:5431:519e:f002:854e:2741:278 [04/Aug/2015:00:01:31 +0000] "GET /1x1.png?u281fd425-s1438646489-i5097.ap.results&zrtd-390.zr4td-1548.zr6td-678.zd-1258.ze-390.zf-971. HTTP/1.1" 200 68 "https://tpc.google syndication.com/sadbundle/7103675352697911246/basic/index.html" "Mozilla/5.0 (Macintosh; Intel Mac OS X 10\_9\_5) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/44.0.2403.125 Safari/537.36" 0.000 https 1438646491.815 Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net


# IPv6 Deployment

## IPv6 Country Deployment for World (XA)



# IPv6 Deployment in Hong Kong

## IPv6 Country Deployment for Hong Kong Special Administrative Region of China (HK)




# IPv6 Deployment in Hong Kong

ASN	AS Name	IPv6 Capable	IPv6 Preferred	Samples
AS9732	SCIG-AS-AP CENTRAL INTERNET SERVICES	94.62%	2.73%	1097
AS4528	HKU-AS-HK The University of Hong Kong	44.85%	40.06%	1233
AS24236	YAHOO-BANGALORE-AS-AP Yahoo Bangalore Network Monitoring Center	40.79%	33.57%	277
AS45187	RACKSPACE-AP Rackspace IT Hosting AS IT Hosting Provider Hong Kong	14.98%	14.54%	1575
AS24334	CYBERPORT-HK-AP Cyberport Hong Kong	4.63%	3.56%	281
AS4158	CITYU-AS-HK City University of Hong Kong	2.84%	2.74%	987
AS38197	SUNHK-DATA-AS-AP Sun Network (Hong Kong) Limited	2.33%	0.39%	258
AS3661	ERX-CUHKNET The Chinese University of Hong Kong	0.70%	0.51%	1573
AS9381	NEWTT-IP-AP Wharf TT Ltd.	0.41%	0.04%	15012
AS7651	LINGNAN-AS-AP Lingnan University	0.35%	0.00%	282
AS24023	JNPR-APAC-AS-AP Juniper Networks HK Ltd.	0.31%	0.08%	1285
AS4760	HKTIMS-AP PCCW Limited	0.29%	0.06%	205465


# IPv6 Deployment in the US

## IPv6 Country Deployment for United States of America (US)







# IPv6 Deployment in Comcast



# Measuring Platforms



# Measuring Browsers



## 2. Measuring DNS Behaviours

# Measuring DNSSEC

Client is given 4 unique URLs to load:

- DNSSEC-validly signed DNS name
- DNSSEC-invalidly signed DNS name
- Unsigned DNS name (control)
- Result reporting URL (10 second timer)

All DNS is IPv4

# What We See (DNS Log)

1438646489.920 [ap] 04-Aug-2015 00:01:29.920 queries: client 202.188.0.254#14118: (Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 157

1438646489.920 [ap] 04-Aug-2015 00:01:29.920 queries: client 202.188.0.254#2911: (04u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: 04u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 145

1438646489.921 [ap] 04-Aug-2015 00:01:29.921 queries: client 202.188.0.254#40461: (Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odu-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 145

1438646489.922 [ap] 04-Aug-2015 00:01:29.922 queries: client 202.188.0.254#48755: (06u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: 06u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 157

1438646489.923 [ap] 04-Aug-2015 00:01:29.923 queries: client 202.188.0.254#12230: (06u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: 06u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 203

1438646489.937 [ap] 04-Aug-2015 00:01:29.937 queries: client 202.188.0.254#11044: (Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 405

1438646489.938 [ap] 04-Aug-2015 00:01:29.938 queries: client 202.188.0.254#58615: (Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Ods-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 417

1438646489.939 [ap] 04-Aug-2015 00:01:29.939 queries: client 202.188.0.254#47094: (Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 405

1438646489.941 [ap] 04-Aug-2015 00:01:29.941 queries: client 202.188.0.254#64994: (Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odi-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 417


1438646490.730 [ap] 04-Aug-2015 00:01:30.730 queries: client 202.188.0.254#42282: (04u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: 04u-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 203

1438646491.466 [ap] 04-Aug-2015 00:01:31.466 queries: client 202.188.0.254#36631: (Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN A -ED () 0 161

1438646491.466 [ap] 04-Aug-2015 00:01:31.466 queries: client 202.188.0.254#52006: (Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net): query: Odu-results-u281fd425-s1438646489-i5097.ap.dotnxdomain.net. IN AAAA -ED () 0 173

# DNSSEC Validation


## Use of DNSSEC Validation for World (XA)






# DNSSEC Validation in Hong Kong

Use of DNSSEC Validation for Hong Kong Special Administrative Region of China (HK)




# DNSSEC Validation in Hong Kong

Use of DNSSEC Validation for Hong Kong Special Administrative Region of China (HK)



# DNSSEC Validation in Sweden

## Use of DNSSEC Validation for Sweden (SE)



# What Else?

## DNSSEC Crypto Support:

How many users who use DNSSEC validating resolvers correctly validate when the signatures use ECDSA (as distinct from RSA)

# Answering the ECC question - DNS + WEB

Data collection: 2/3/15 – 19/3/15

1,830,668 clients who appear to be exclusively using RSA DNSSEC-Validating resolvers

## ECC Results:

Success: 79.9% 1,461,772 Saw fetches of the ECC DNSSEC RRs and the well-signed named URL, but not the badly signed named URL

## Failure (fetched both URLs):

Mixed Resolvers	5.1%	93,746	Used both ECDSA-Validating and non-validating resolvers
NO ECC	13.3%	243,794	Saw A, DS, no DNSKEY, fetched both URLs
Mixed	1.3%	24,420	Saw some DNSSEC queries, fetched both URLs
No Validation	0.4%	6,836	Did not fetch any DNSSEC RRs

**Apparent Fail: 20.1% 368,796**

1 in 5 clients that use resolvers that perform DNSSEC validation with RSA fail to validate with ECDSA

# ECC Results

- These results show that 80% of clients who appeared to use RSA DNSSEC-Validating resolvers were also seen to perform validation using ECDSA
- Two thirds of the the remaining clients fetched both objects (13% of the total), but did not fetch any DNSKEY RRs.
- Of the remainder (5%), most were using a validating resolver (which returned SERVFAIL for the badly signed object), and then the client failed over to a non-validating resolver \*

\* This is curious, because these clients did not failover to a non-validating resolver on a badly signed RSA structure

# What Else?

- The “market” for DNS resolution: how many users send their queries through Google’s Public DNS servers?
- How many users use resolvers located in a foreign country?
- Which countries?

# Foreign (CC) Resolution: Top Resolvers by AS

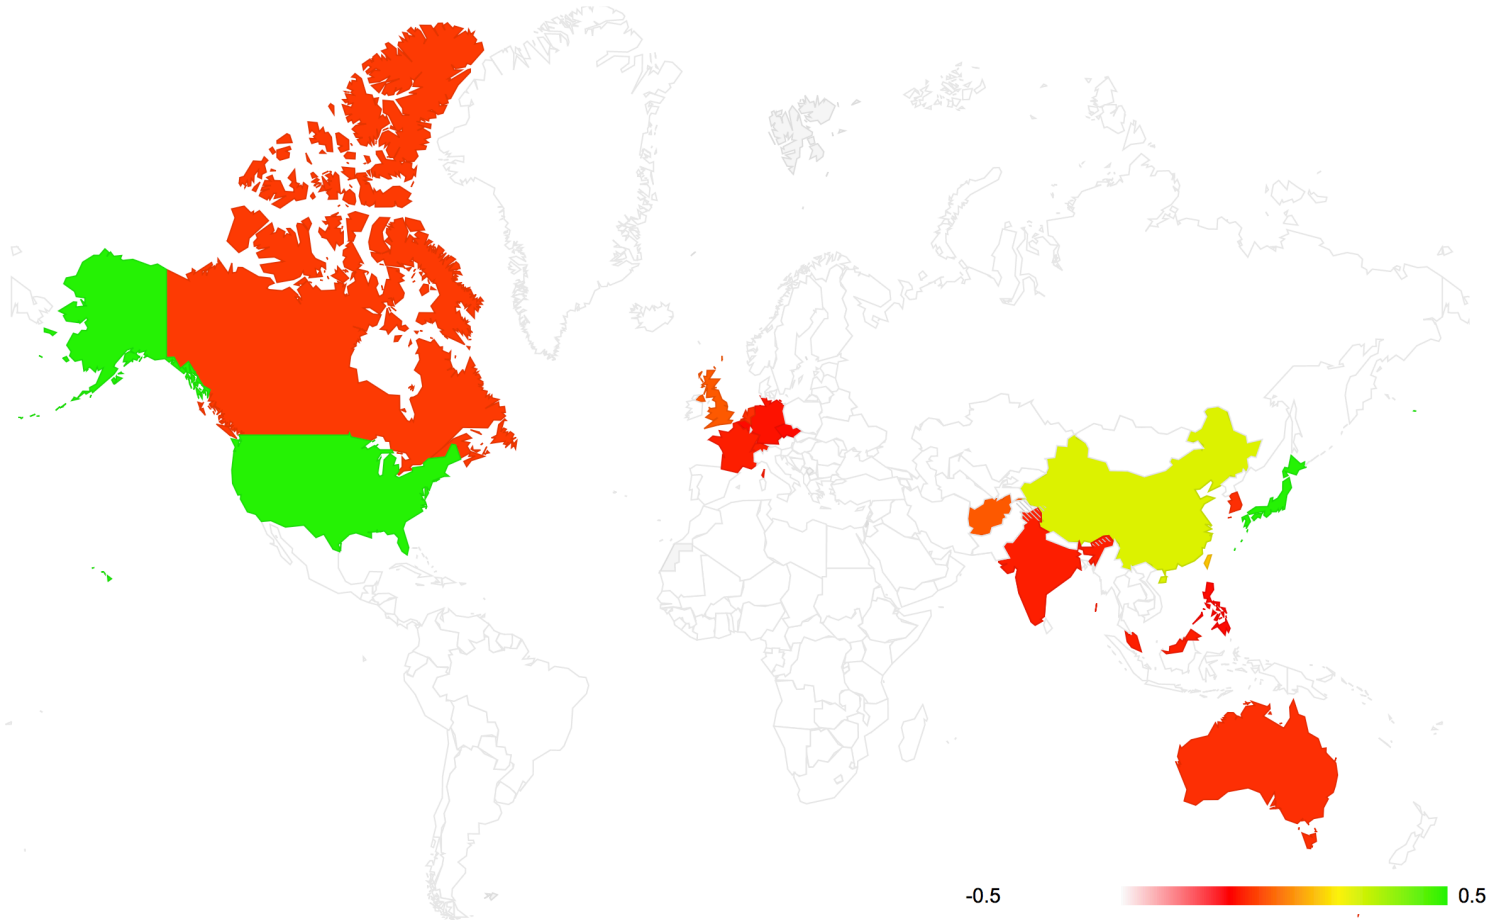
Rank	AS	Use	AS Name
1	15169	42.69%	GOOGLE - Google Inc.,US
2	3356	7.47%	LEVEL3 - Level 3 Communications, Inc.,US
3	36692	7.05%	OPENDNS - OpenDNS, LLC,US
4	19994	2.56%	RACKSPACE - Rackspace Hosting,US
5	174	1.87%	COGENT-174 - Cogent Communications,US
6	16880	1.70%	AS2-TRENDMICRO-COM - TREND MICRO INCORPORATED,US
7	2914	1.09%	NTT-COMMUNICATIONS-2914 - NTT America, Inc.,US
8	4134	0.91%	CHINANET-BACKBONE No.31,Jin-rong Street,CN
9	29791	0.70%	VOXEL-DOT-NET - Voxel Dot Net, Inc.,US
10	3462	0.67%	HINET Data Communication Business Group,TW
11	9121	0.64%	TTNET Turk Telekomunikasyon Anonim Sirketi,TR
12	3303	0.64%	SWISSCOM Swisscom (Switzerland) Ltd,CH
13	6939	0.63%	HURRICANE - Hurricane Electric, Inc.,US
14	6147	0.50%	Telefonica del Peru S.A.A.,PE
15	6713	0.48%	IAM-AS,MA
16	8048	0.47%	CANTV Servicios, Venezuela,VE
17	3257	0.47%	TINET-BACKBONE Tinet SpA,DE
18	13238	0.43%	YANDEX Yandex LLC,RU
19	45595	0.41%	PKTELECOM-AS-PK Pakistan Telecom Company Limited,PK
20	9299	0.40%	IPG-AS-AP Philippine Long Distance Telephone Company,PH
21	7643	0.39%	VNPT-AS-VN Vietnam Posts and Telecommunications (VNPT),VN
22	45758	0.39%	TRIPLETNET-AS-AP Triplet Internet Internet service provider Bangkok,TH
23	8151	0.38%	Uninet S.A. de C.V.,MX
24	7470	0.35%	TRUEINTERNET-AS-AP TRUE INTERNET Co.,Ltd.,TH
25	4837	0.35%	CHINA169-BACKBONE CNCGROUP China169 Backbone,CN

Total: 21,770,772 (28% of total) end user query sets



# Offshore DNS from HK Users

Foreign Resolver Distribution for Hong Kong Special Administrative Region of China (2014) (57610 samples, 7174 foreign resolution instances (12%))



# Offshore DNS from HK Users


Foreign Resolver Distribution for Hong Kong Special Administrative Region of China (2014) (57610 samples, 7174 foreign resolution instances (12%)) **DNS Resolution Distribution for Hong Kong Special Administrative Region of China**

CC	Country	Resolver Samples	Resolver Share
HK	Hong Kong Special Administrative Region of China	50436	87.55%
XX	Total Foreign DNS resolution	7174	12.45%
ZZ	Google Public DNS	5025	8.72%
<a href="#">US</a>	United States of America	1042	1.81%
<a href="#">JP</a>	Japan	314	0.55%
<a href="#">CN</a>	China	166	0.29%
<a href="#">SG</a>	Singapore	156	0.27%
<a href="#">TW</a>	Taiwan	116	0.20%
<a href="#">AF</a>	Afghanistan	54	0.09%
<a href="#">GB</a>	United Kingdom of Great Britain and Northern Ireland	52	0.09%
<a href="#">CA</a>	Canada	35	0.06%
<a href="#">KR</a>	Republic of Korea	31	0.05%
<a href="#">AU</a>	Australia	28	0.05%
<a href="#">NL</a>	Netherlands	25	0.05%
<a href="#">MY</a>	Malaysia	18	0.03%
<a href="#">FR</a>	France	17	0.03%
<a href="#">IN</a>	India	17	0.03%
<a href="#">EU</a>	European Union	16	0.03%
<a href="#">CH</a>	Switzerland	15	0.03%
<a href="#">DE</a>	Germany	14	0.02%
<a href="#">BE</a>	Belgium	8	0.01%
<a href="#">PH</a>	Philippines	7	0.01%
<a href="#">CZ</a>	Czech Republic	7	0.01%



# Market Penetration of Google's Public DNS

## Use of DNSSEC Validation for World (XA)



# Why is this happening?

- ❑ It's Google: Google's Public DNS (all instances are mapped to the US in this per-AS analysis)
- ❑ Users' efforts to circumvent DNS-based geo-loc content access controls (think Netflix!)
- ❑ 3<sup>rd</sup> party DNS query monitoring/stalking (yes, there is some of this going on, but that's a talk for another time!)
- ❑ Virus contamination of the host (yes, captured systems often show a redirected DNS config)
- ❑ <insert your favourite theory here>

# 3. Digital Stalking

Who's Watching?



# Some Stalker Numbers

In the first 248 days of 2014 we saw:

- 123,110,633 unique end-user IP addresses presented to our servers from these test scripts
- 317,309 of these end-user IP addresses presented HTTP GET strings to us that were subsequently presented to us from a different client IP address!

That's some **1 in 400\*** users that seem to have attracted some kind of digital stalker!

\* Or maybe a bit more, due to NATs hiding multiple end users behind a single public IP address






Take the tour to see  
what's new »



# Committed to you, your privacy and an open Web

## Keep your Firefox in Sync

Access your bookmarks, passwords  
and more from any device.



Get started with Sync

*Create an account from the menu panel*



# Online Privacy? Really?

It's hard to believe that today's Internet respects personal privacy when it seems that around 1 in 400 users have attracted some kind of digital stalker that tracks the URLs they visit.


# Stalking Rates by Country

CC	Samples	stalked	Rate/1,000,000	Country
IR	674	111	164,688	Iran (Islamic Republic of)
LA	28,506	2,875	100,855	Lao People's Democratic Republic
MO	38,761	2,954	76,210	Macao Special Administrative Region of China
SG	240,188	17,406	72,468	Singapore
HK	486,101	22,136	45,537	Hong Kong Special Administrative Region of China
CN	10,419,638	435,040	41,751	China
GB	872,124	28,845	33,074	United Kingdom of Great Britain and Northern Ireland
TW	1,769,367	36,823	20,811	Taiwan
JP	1,500,779	23,971	15,972	Japan
AU	293,193	4,620	15,757	Australia
US	4,491,711	53,370	11,881	United States of America
MY	1,035,434	10,214	9,864	Malaysia
AL	437,399	4,043	9,243	Albania
CA	947,922	6,244	6,587	Canada
KH	143,886	897	6,234	Cambodia
MM	16,411	97	5,910	Myanmar
MK	458,820	2,214	4,825	The former Yugoslav Republic of Macedonia
BZ	8,139	35	4,300	Belize
MN	57,622	233	4,043	Mongolia
NZ	344,951	1,385	4,015	New Zealand
CV	3,742	14	3,741	Cape Verde
ME	223,005	775	3,475	Montenegro
FJ	14,892	47	3,156	Fiji
SR	44,116	136	3,082	Suriname
AW	11,123	34	3,056	Aruba

The top 25 countries in terms of observed URL stalking rates

# Stalking Delay

Distribution of Stalking Delay



# Top 25 International Stalkers

Rank	IP Net	#	AVG Delay	AS	Description
1	119.147.146.0	205,033	130.7	4134	CHINANET-BACKBONE No.31,Jin-rong Street,CN
2	101.226.33.0	6,198	1,576.1	4812	CHINANET-SH-AP China Telecom (Group),CN
3	180.153.206.0	6,120	1,608.3	4812	CHINANET-SH-AP China Telecom (Group),CN
4	180.153.214.0	3,827	1,561.0	4812	CHINANET-SH-AP China Telecom (Group),CN
5	112.64.235.0	3,819	1,544.9	17621	CNCGROUP-SH China Unicom Shanghai network,CN
6	101.226.66.0	3,603	1,577.3	4812	CHINANET-SH-AP China Telecom (Group),CN
7	180.153.163.0	2,742	1,540.1	4812	CHINANET-SH-AP China Telecom (Group),CN
8	223.27.200.0	2,740	1.8	45796	BBCONNECT-TH-AS-AP BB Connect Co., Ltd.,TH
9	101.226.89.0	2,658	2,230.2	4812	CHINANET-SH-AP China Telecom (Group),CN
10	180.153.201.0	2,628	1,549.4	4812	CHINANET-SH-AP China Telecom (Group),CN
11	101.226.65.0	1,528	1,573.3	4812	CHINANET-SH-AP China Telecom (Group),CN
12	69.41.14.0	1,243	1,127.4	47018	CE-BGPAC - Covenant Eyes, Inc.,US
13	101.226.51.0	1,195	1,627.6	4812	CHINANET-SH-AP China Telecom (Group),CN
14	112.65.193.0	1,038	1,623.9	17621	CNCGROUP-SH China Unicom Shanghai network,CN
15	64.124.98.0	906	1,288.9	6461	ABOVENET - Abovenet Communications, Inc,US
16	180.153.114.0	819	1,632.6	4812	CHINANET-SH-AP China Telecom (Group),CN
17	180.153.205.0	765	1,497.7	4812	CHINANET-SH-AP China Telecom (Group),CN
18	208.184.77.0	649	1,419.5	6461	ABOVENET - Abovenet Communications, Inc,US
19	222.73.77.0	535	1,373.8	4812	CHINANET-SH-AP China Telecom (Group),CN
20	180.153.211.0	517	1,450.6	4812	CHINANET-SH-AP China Telecom (Group),CN
21	180.153.161.0	504	1,675.7	4812	CHINANET-SH-AP China Telecom (Group),CN
22	183.60.153.0	262	451.3	4134	CHINANET-BACKBONE No.31,Jin-rong Street,CN
23	222.73.76.0	255	1,512.7	4812	CHINANET-SH-AP China Telecom (Group),CN
24	101.226.102.0	235	2,012.7	4812	CHINANET-SH-AP China Telecom (Group),CN
25	208.80.194.0	227	10,731.5	13448	WEBSense - Websense, Inc,US

# The Leakiest Browser!



Wow! "Public Security Equipment 110 No 0000000025!"

# 4. Access ISP Market Share

<http://stats.labs.apnic.net/aspop>

# Market Share in HK

Rank	ASN	AS Name	CC	Users (est.)	% of country	% of Internet	Samples
1	AS4760	HKTIMS-AP PCCW Limited	<a href="#">HK</a>	1758644	32.19	0.0569	249956
2	AS9269	HKBN-AS-AP Hong Kong Broadband Network Ltd.	<a href="#">HK</a>	1632070	29.87	0.0528	231966
3	AS9304	HUTCHISON-AS-AP Hutchison Global Communications	<a href="#">HK</a>	483944	8.86	0.0157	68783
4	AS9908	HKCABLE2-HK-AP HK Cable TV Ltd	<a href="#">HK</a>	274826	5.03	0.0089	39061
5	AS9381	NEWTT-IP-AP Wharf TT Ltd.	<a href="#">HK</a>	137022	2.51	0.0044	19475
6	AS10026	PACNET Pacnet Global Ltd	<a href="#">HK</a>	118525	2.17	0.0038	16846
7	AS4515	ERX-STAR PCCW IMSBiz	<a href="#">HK</a>	98740	1.81	0.0032	14034
8	AS38819	HKCSL-AS-AP HKCSL GPRS NETWORK	<a href="#">HK</a>	91662	1.68	0.003	13028
9	AS55536	PSWITCH-HK PACSWITCH GLOBAL IP NETWORK	<a href="#">HK</a>	89326	1.63	0.0029	12696
10	AS18116	HGC-AS-AP Hutchison Global Crossing	<a href="#">HK</a>	87363	1.6	0.0028	12417
11	AS9231	IPEOPLESNET-AS-AP China Mobile Hong Kong Company Limited	<a href="#">HK</a>	80137	1.47	0.0026	11390
12	AS17924	SMARTONE-MB-AS-AP SmarTone Mobile Communications Ltd	<a href="#">HK</a>	60177	1.1	0.0019	8553
13	AS9474	SMARTONE-AS-AP SmarTone Telecommunications Ltd.	<a href="#">HK</a>	46541	0.85	0.0015	6615
14	AS10103	HKBN-AS-AP HK Broadband Network Ltd.	<a href="#">HK</a>	37796	0.69	0.0012	5372
15	AS10118	HTCL-IAS-HK-AP Hutchison Telephone Company Limited	<a href="#">HK</a>	36586	0.67	0.0012	5200
16	AS9444	HKT-AS-AP Hong Kong Telecommunications (HKT) Limited	<a href="#">HK</a>	29501	0.54	0.001	4193
17	AS4637	ASN-TELSTRA-GLOBAL Telstra Global	<a href="#">HK</a>	21128	0.39	0.0007	3003
18	AS9229	SPEEDCAST-AP SPEEDCAST Limited	<a href="#">HK</a>	17631	0.32	0.0006	2506
19	AS3661	ERX-CUHKNET The Chinese University of Hong Kong	<a href="#">HK</a>	15098	0.28	0.0005	2146
20	AS9221	HSBC-HK-AS HSBC HongKong	<a href="#">HK</a>	14768	0.27	0.0005	2099

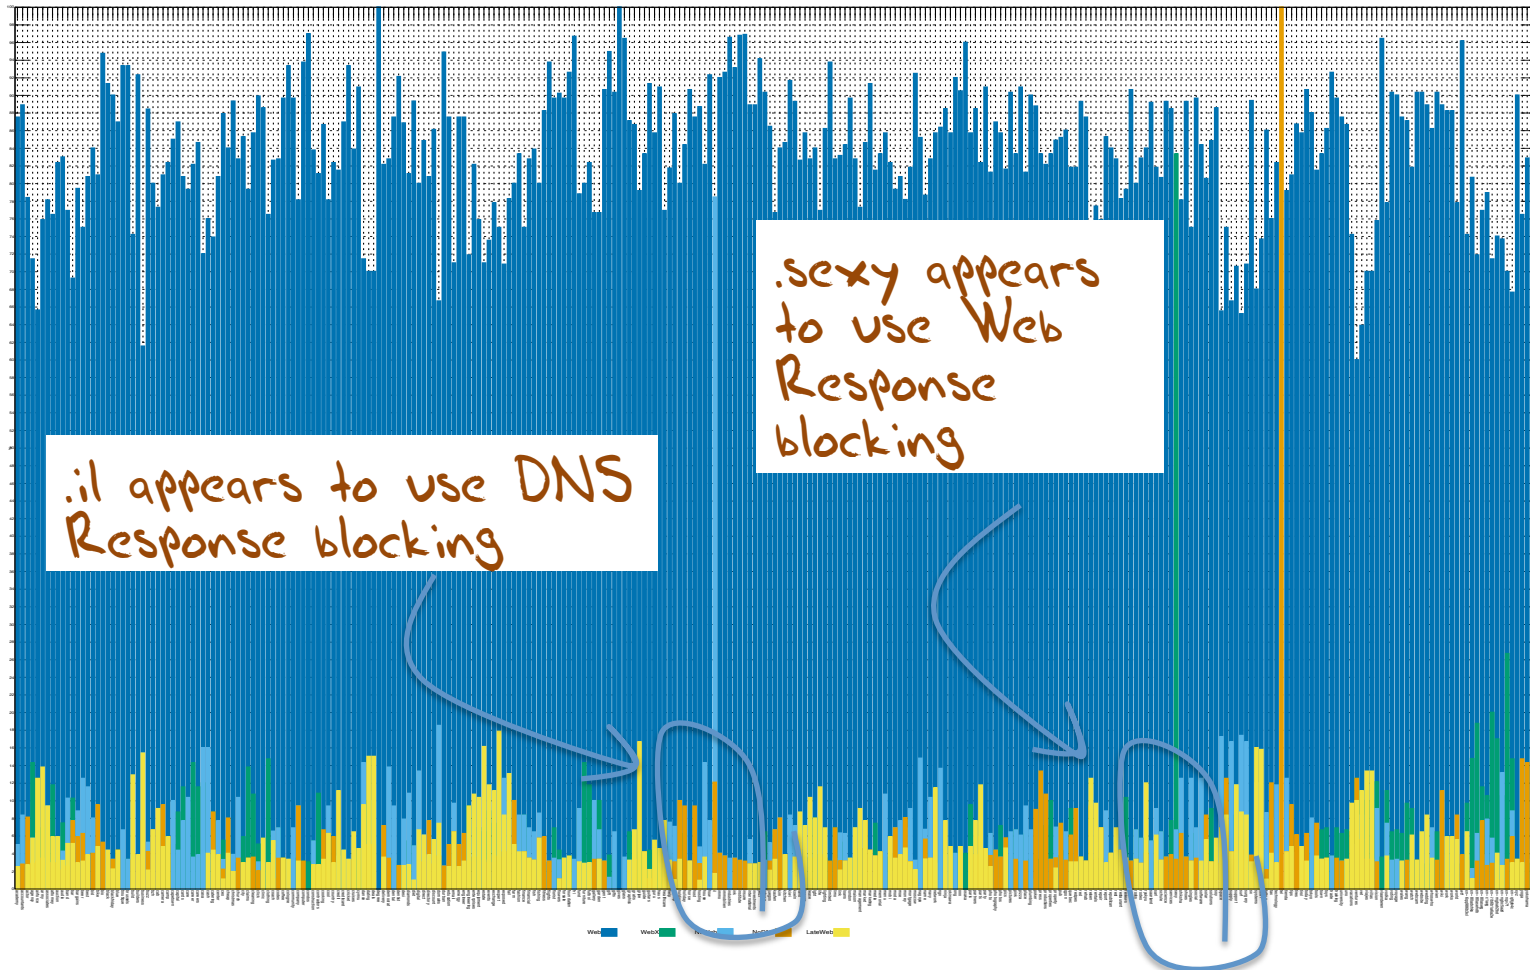
# What Else?

Analysis of failure patterns to detect evidence of structured interception of DNS and Web retrieval



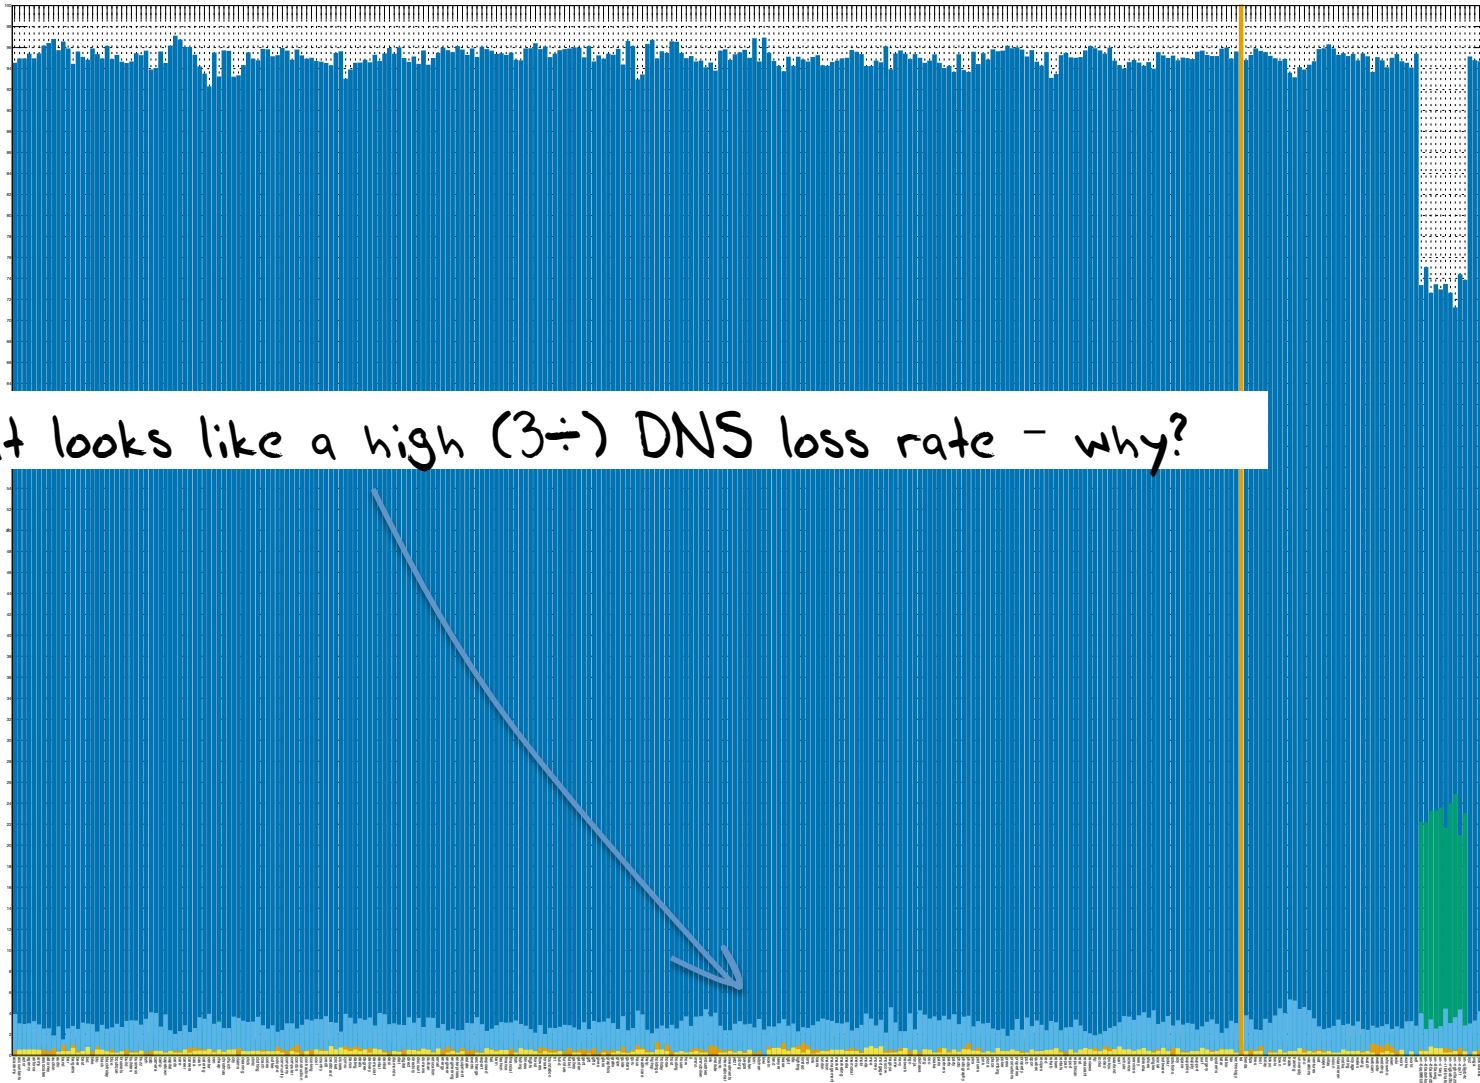
# Content Blocking in Iran?

Iran (Islamic Republic of) (IR) - 11025 Measurements



# Hong Kong

Hong Kong Special Administrative Region of China (HK) - 248750 Measurements



That looks like a high (3÷) DNS loss rate - why?

# What Else?

- This approach allows us to analyze user behaviour when presented with particular tests
  - DNS: response size, TCP behaviour, resolver distribution, matching resolvers to users, resolver timers, EDNS0 use, EDNS0 client subnet use and accuracy, dual stack behaviour, response size,...
  - Web: Protocol preference, dual stack behaviour, response size, fragmentation behaviour, ...

# But...

- It's not a general purpose compute platform, so it can't do many things
  - Ping, traceroute, etc
  - Send data to any destination
  - Pull data from any destination
  - Use different protocols
- This is a “many-to-one” styled setup where the server instrumentation provides insight on the inferred behaviour of the edges

# In Summary...

- Measuring what happens at the user level by measuring some artifact or behaviour in the infrastructure and inferring some form of user behaviour is always going to be a guess of some form
- If you really want to measure user behaviour then its useful to trigger the user to behave in the way you want to study or measure
- The technique of embedding simple test code behind ads is one way of achieving this objective
  - for certain kinds of behaviours relating to the DNS and to URL fetching

Thanks to the folk at Google Research for their generous support of our work!

APNIC Labs:

Geoff Huston

George Michaelson

research@apnic.net

Questions?