Measuring the End User

Geoff Huston
APNIC Labs, May 2015


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?

etc

"Measurable" Questions?

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

…

Users vs Infrastructure

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

Private Data

- 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 the entire Internet that do 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 endpoints devices on today's Internet are capable of making IPv6 connections?

How to measure a million end devices for their IPv6 capability?

How to measure a million end devices for their IPv6 capability?

a) Be

How to measure a million end devices for their IPv6 capability?

a) Be Google

OR

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

Ads are ubiquitous

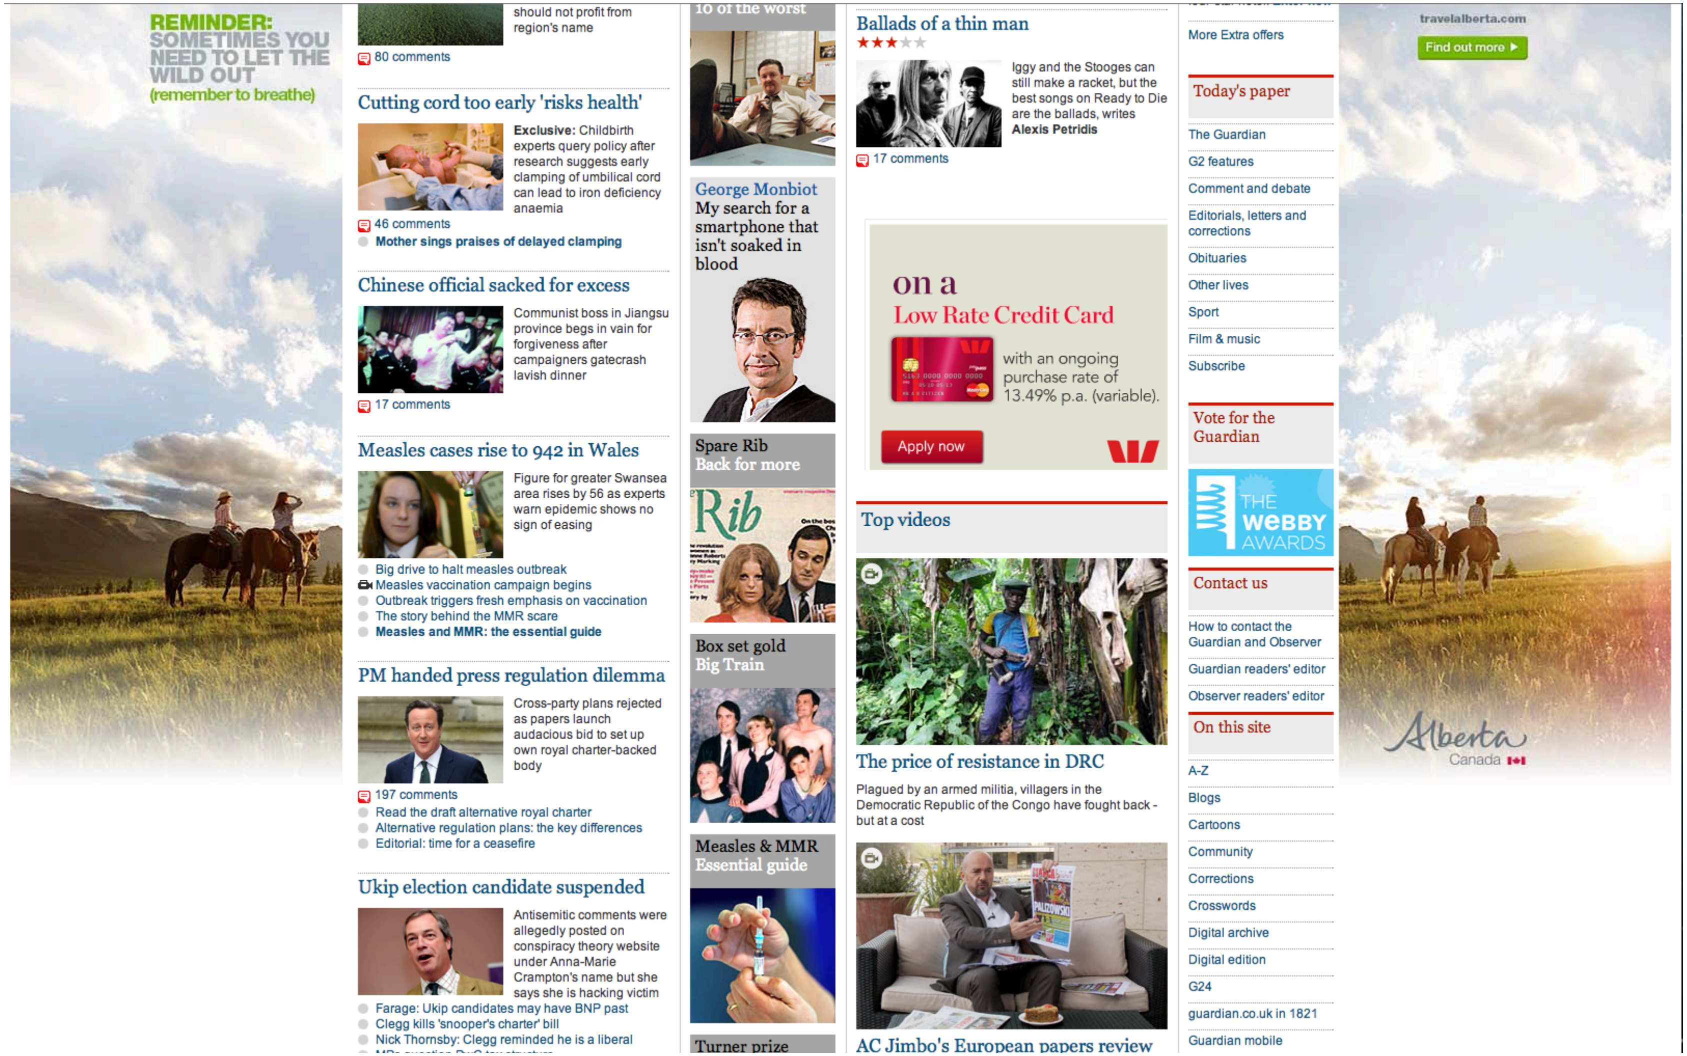

Ads are ubiquitous

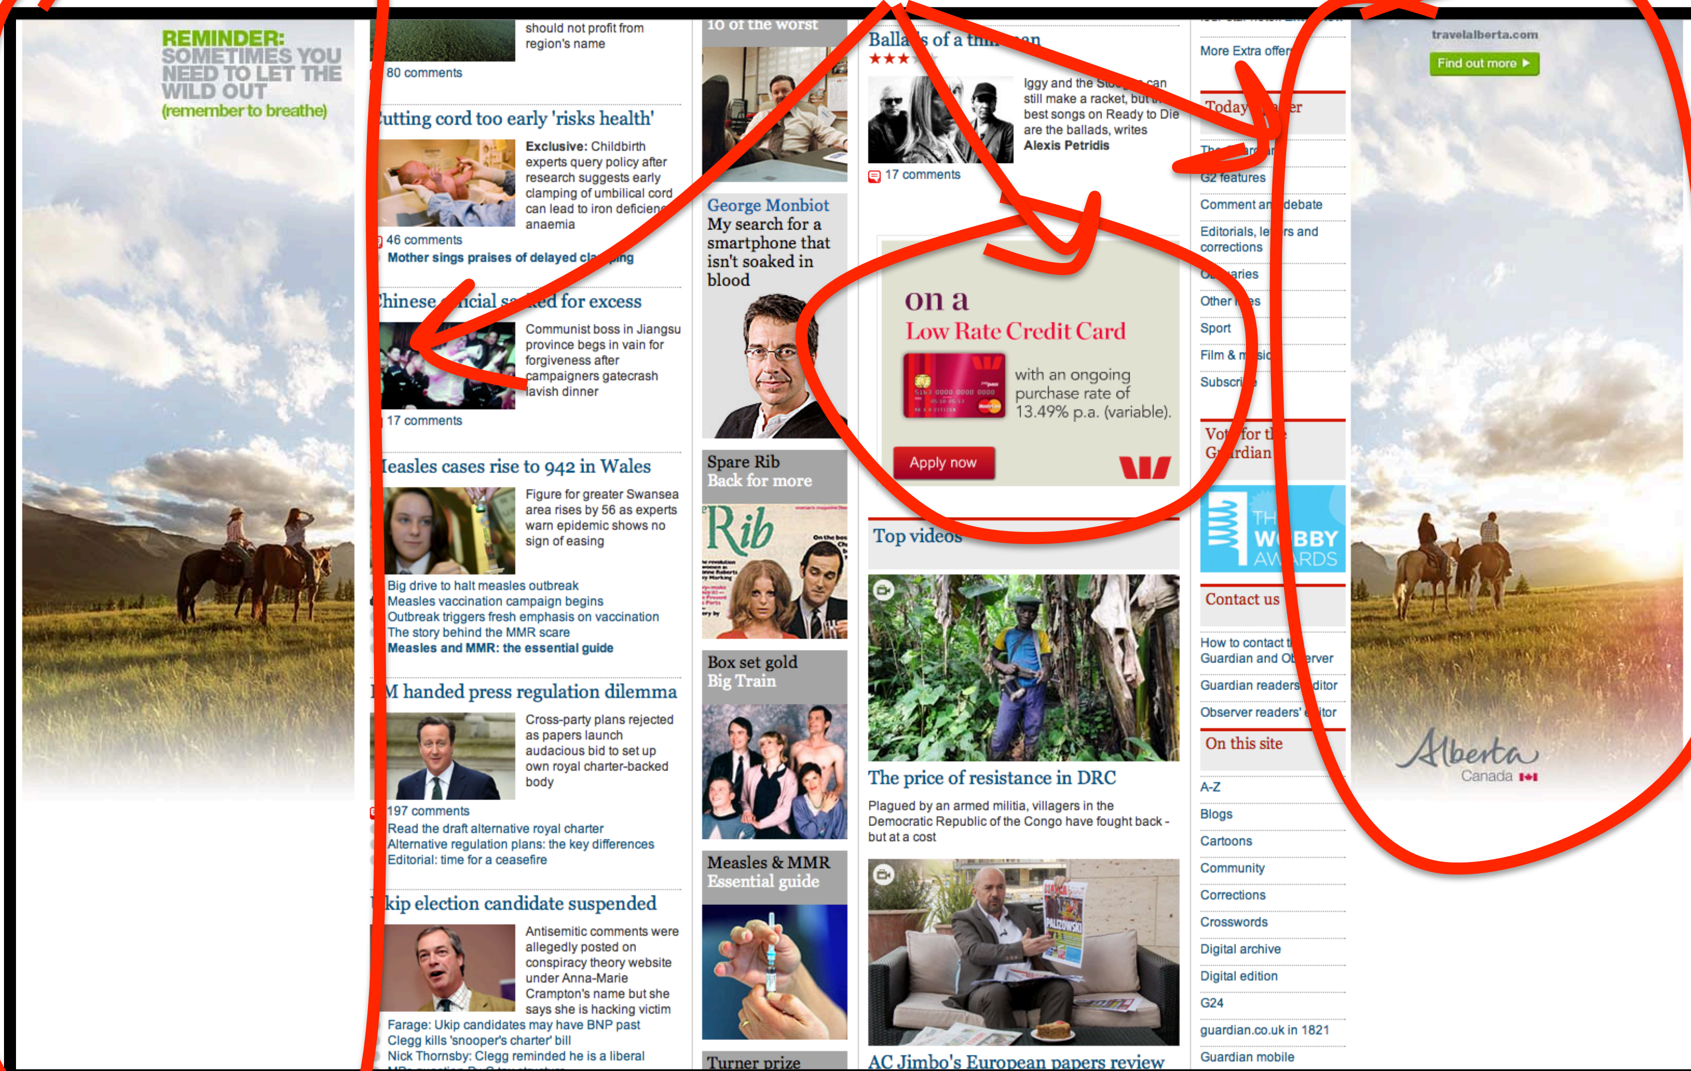

Ads are ubiquitous

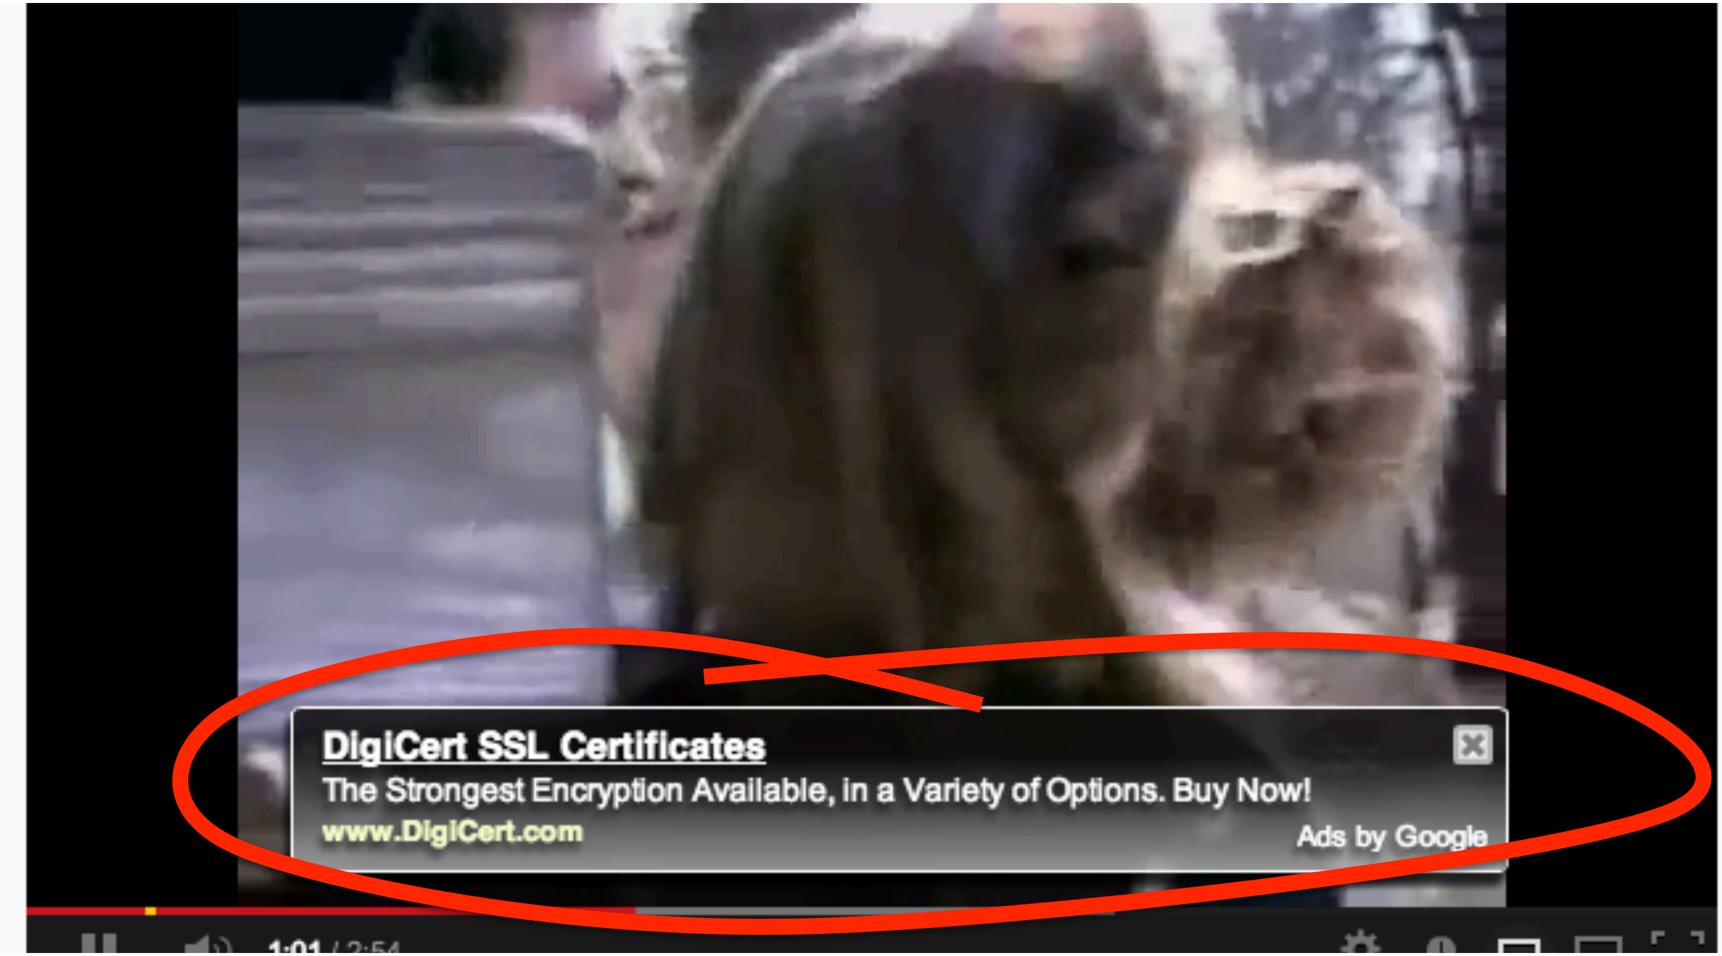

Ads use active scripts

- 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

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

APNIC's measurement technique

- Craft Flash/Actionscript which fetches network assets to measure.
- Assets are reduced to a notional ‘1x1’ image which is not added to the DOM and is not displayed
- Assets can be named (DNS resolution via local gethostbyname() styled API within the browser’s Flash engine) or use literals (bypass DNS resolution)
- Encode data transfer in the name of fetched assets
  – 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
  – Or just use the server logs!

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 24h algorithm tries hard to even placement but in the end, will 'soak' place your ad to achieve enough views, to bill you $100

Ad Placement Training – Day 1

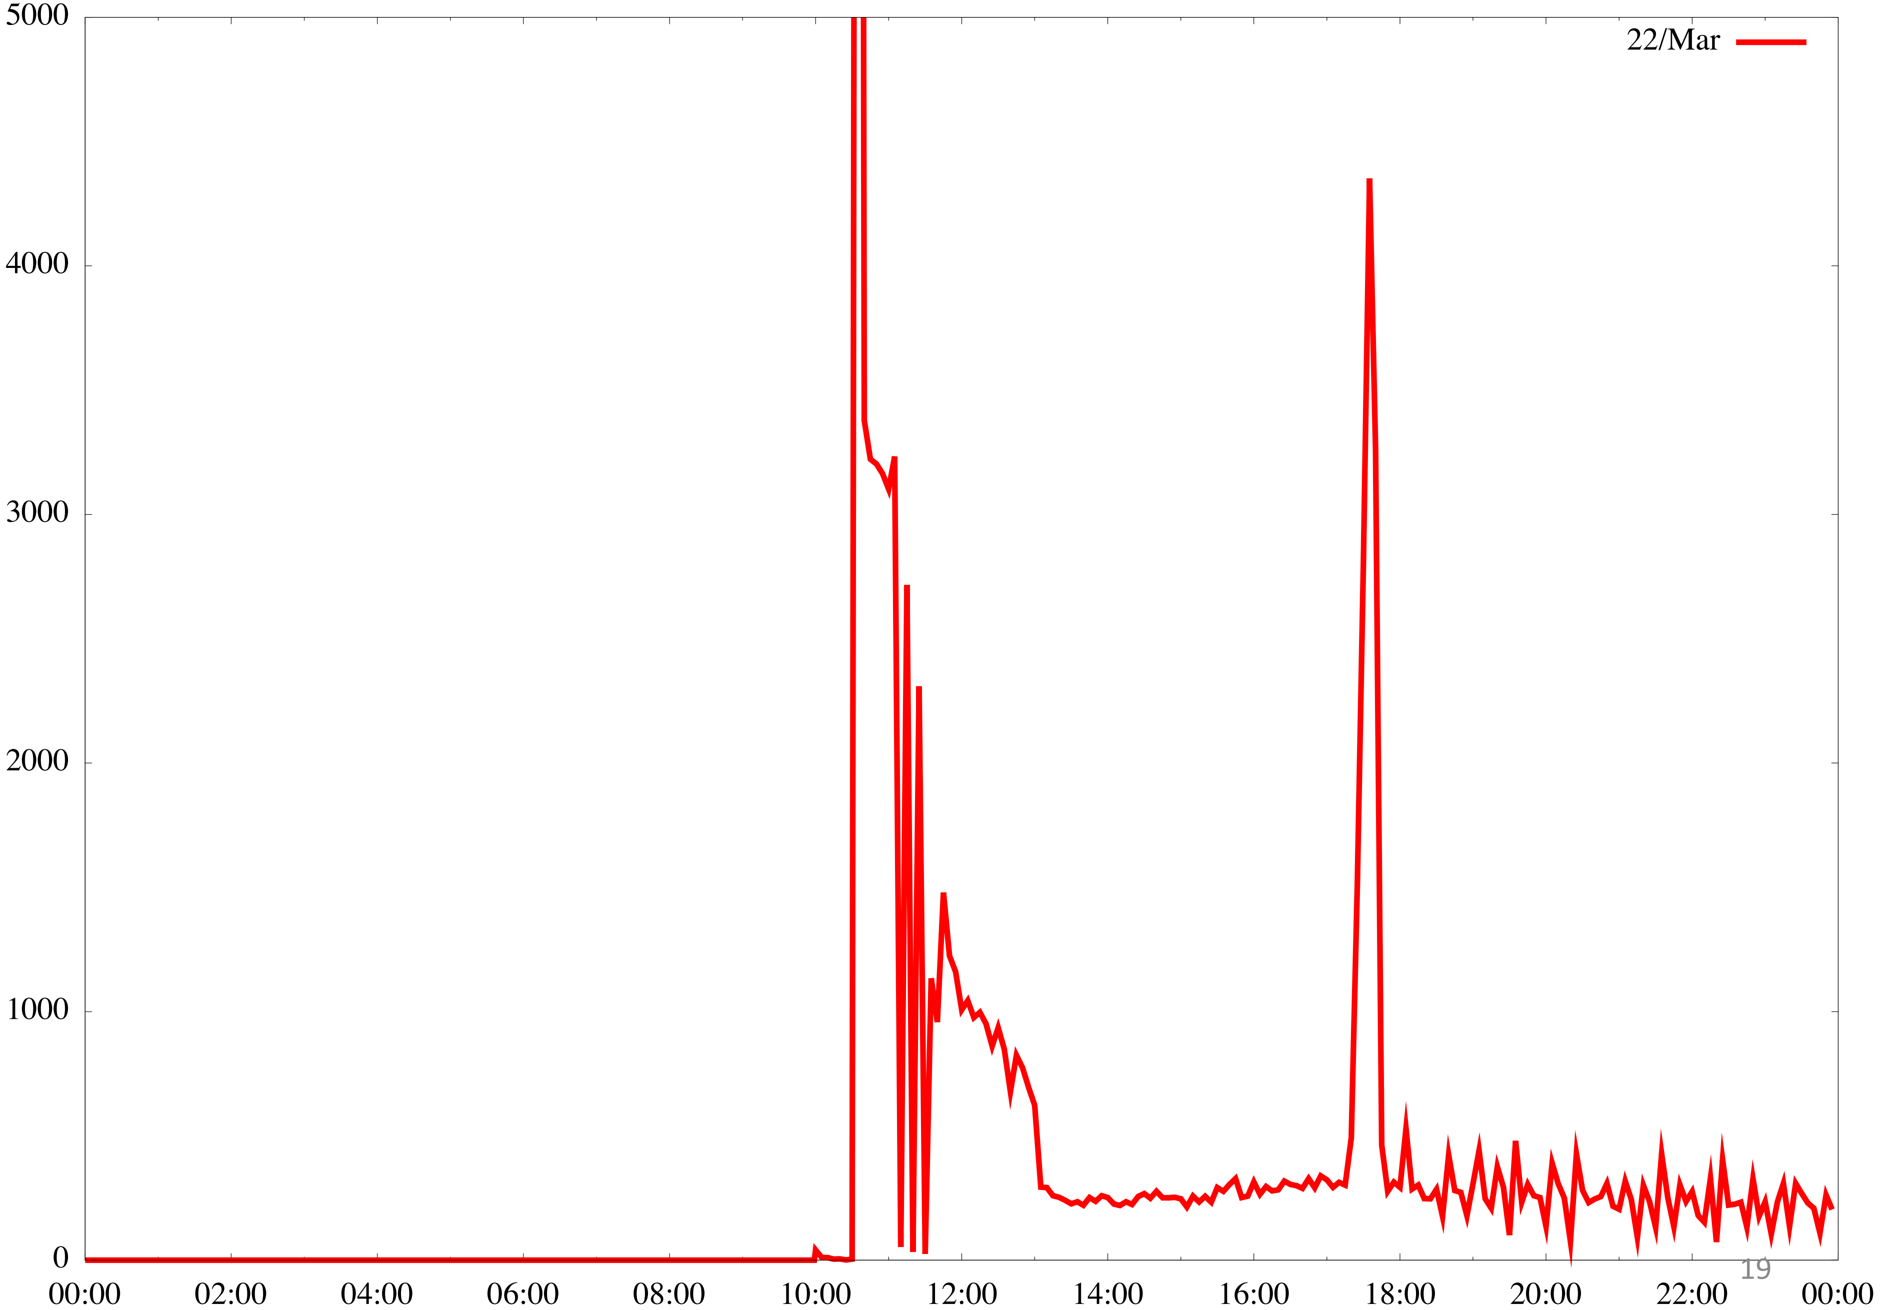

Ad Placement Training – Day 2

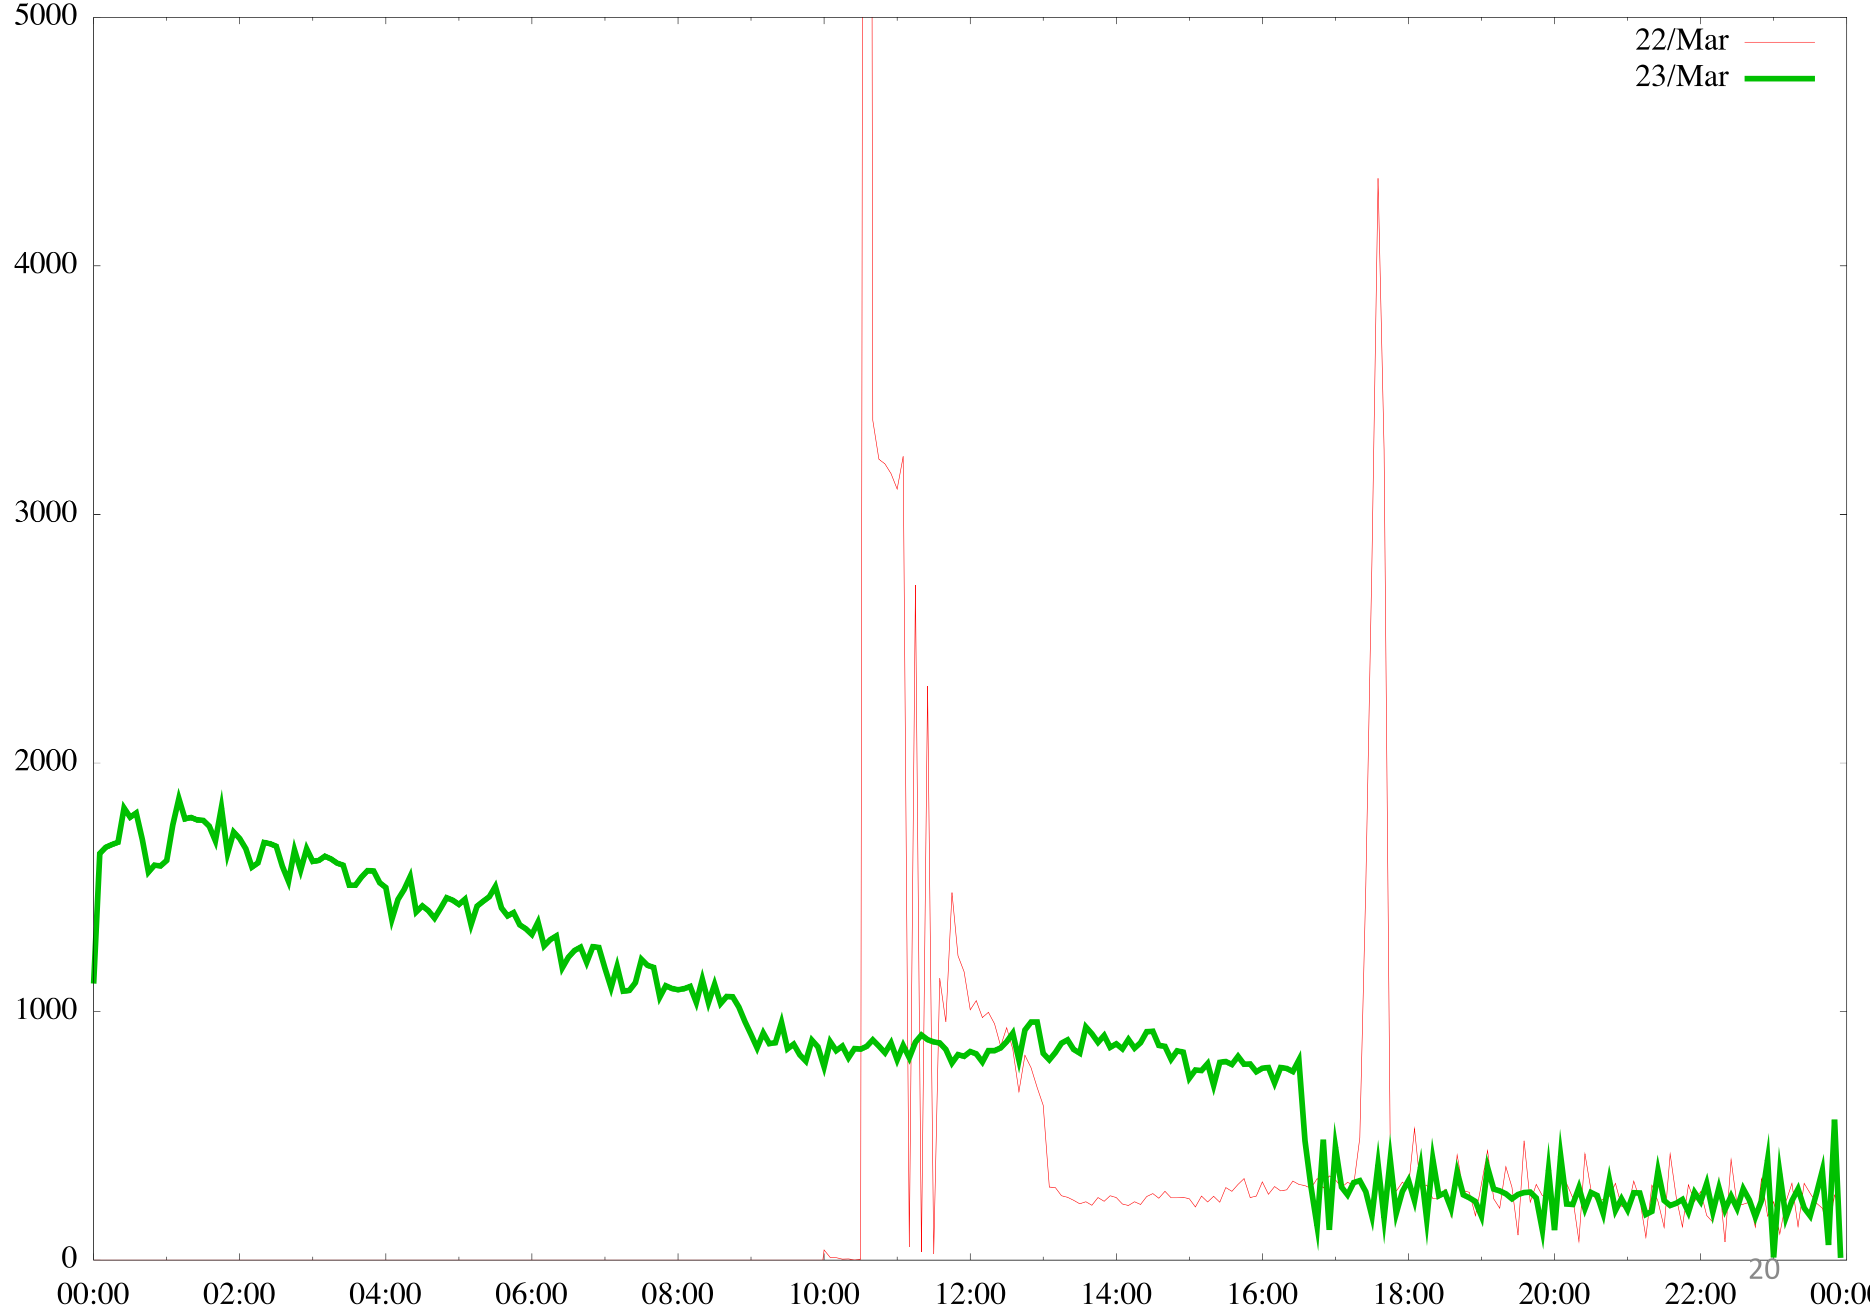

Ad Placement Training – Day 3

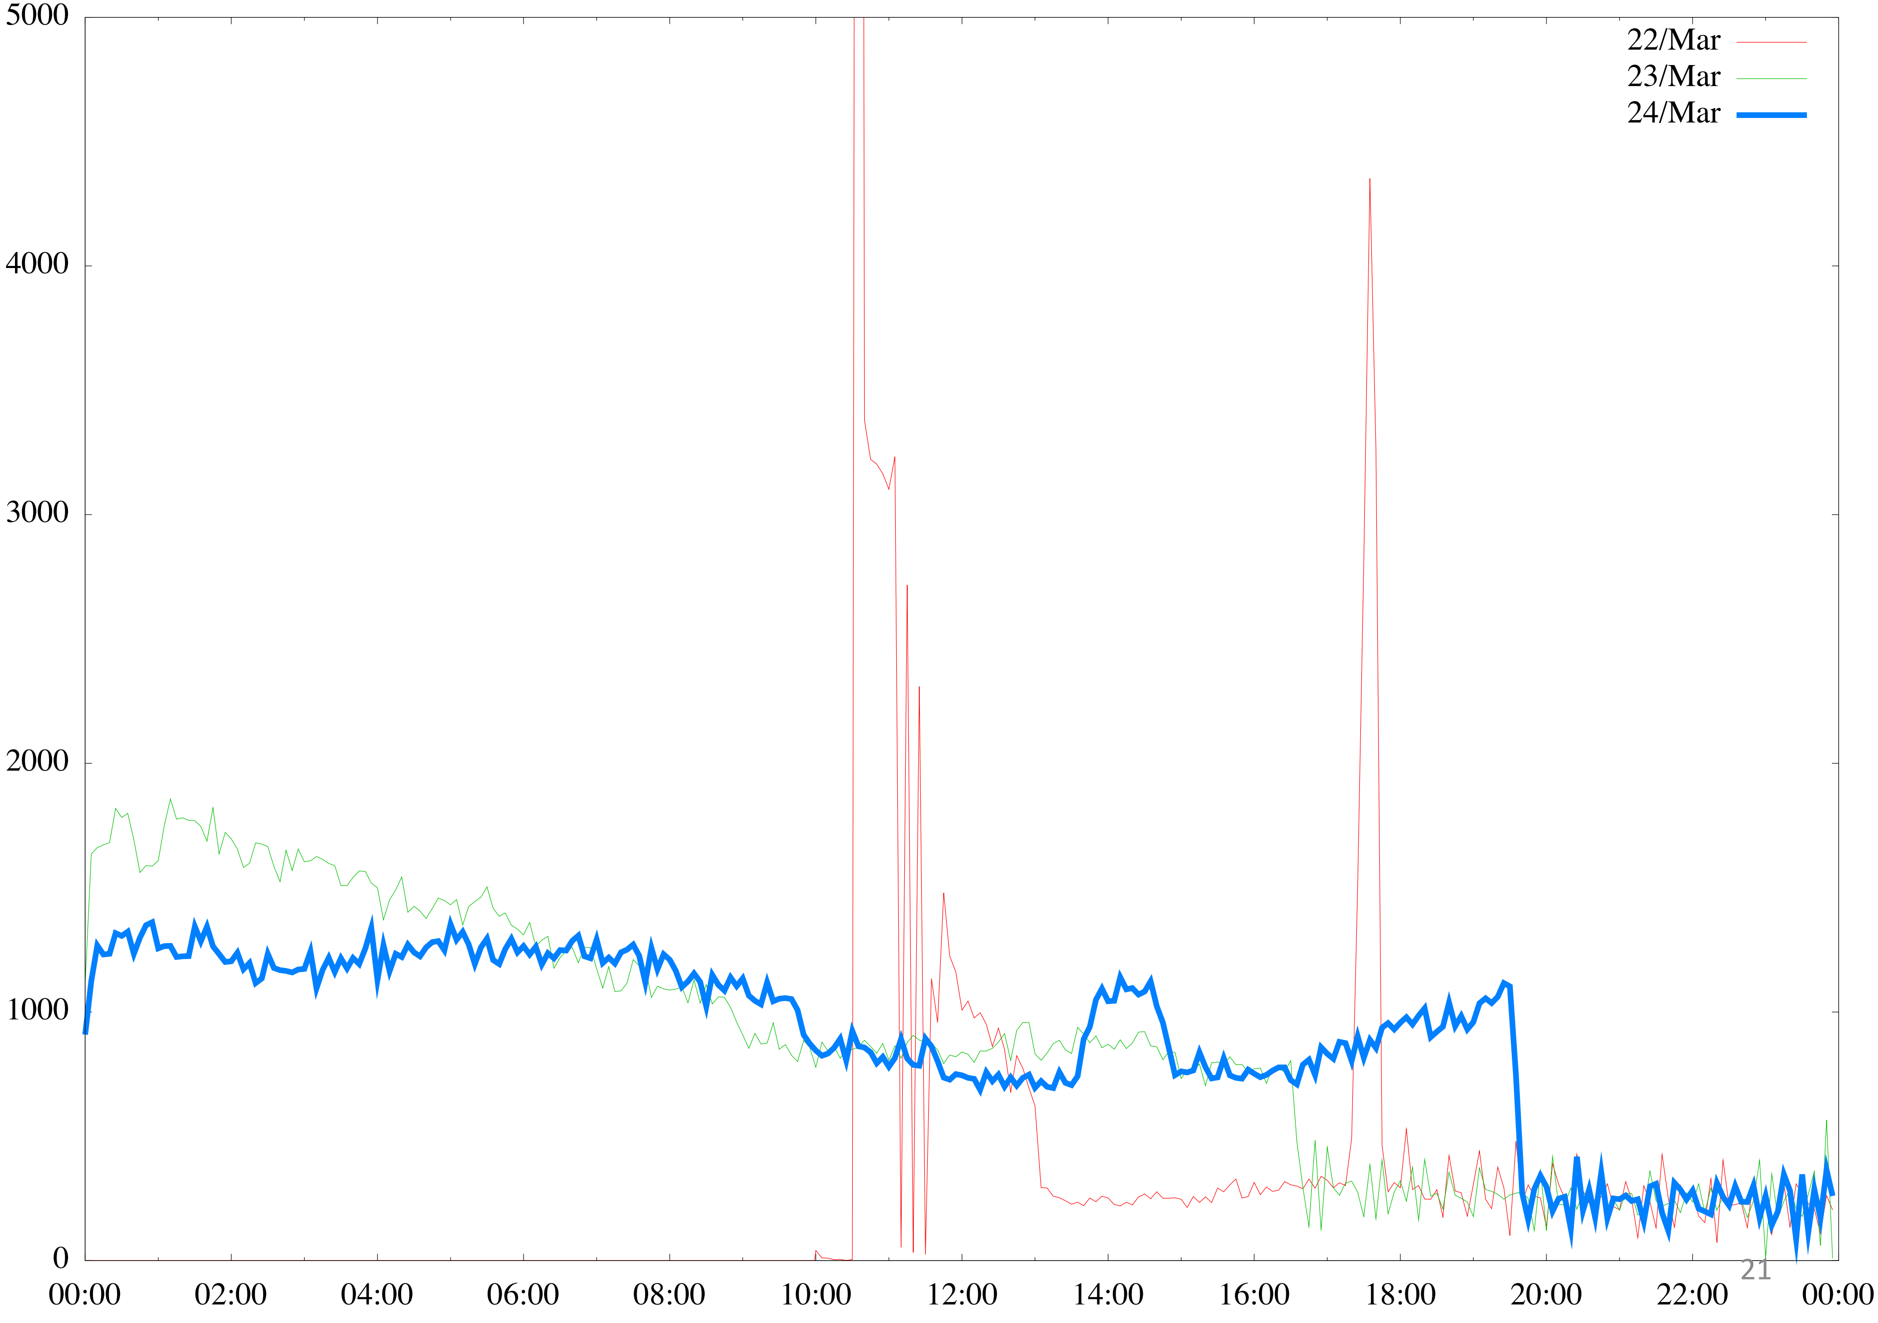

Ad Placement Training – Day 4

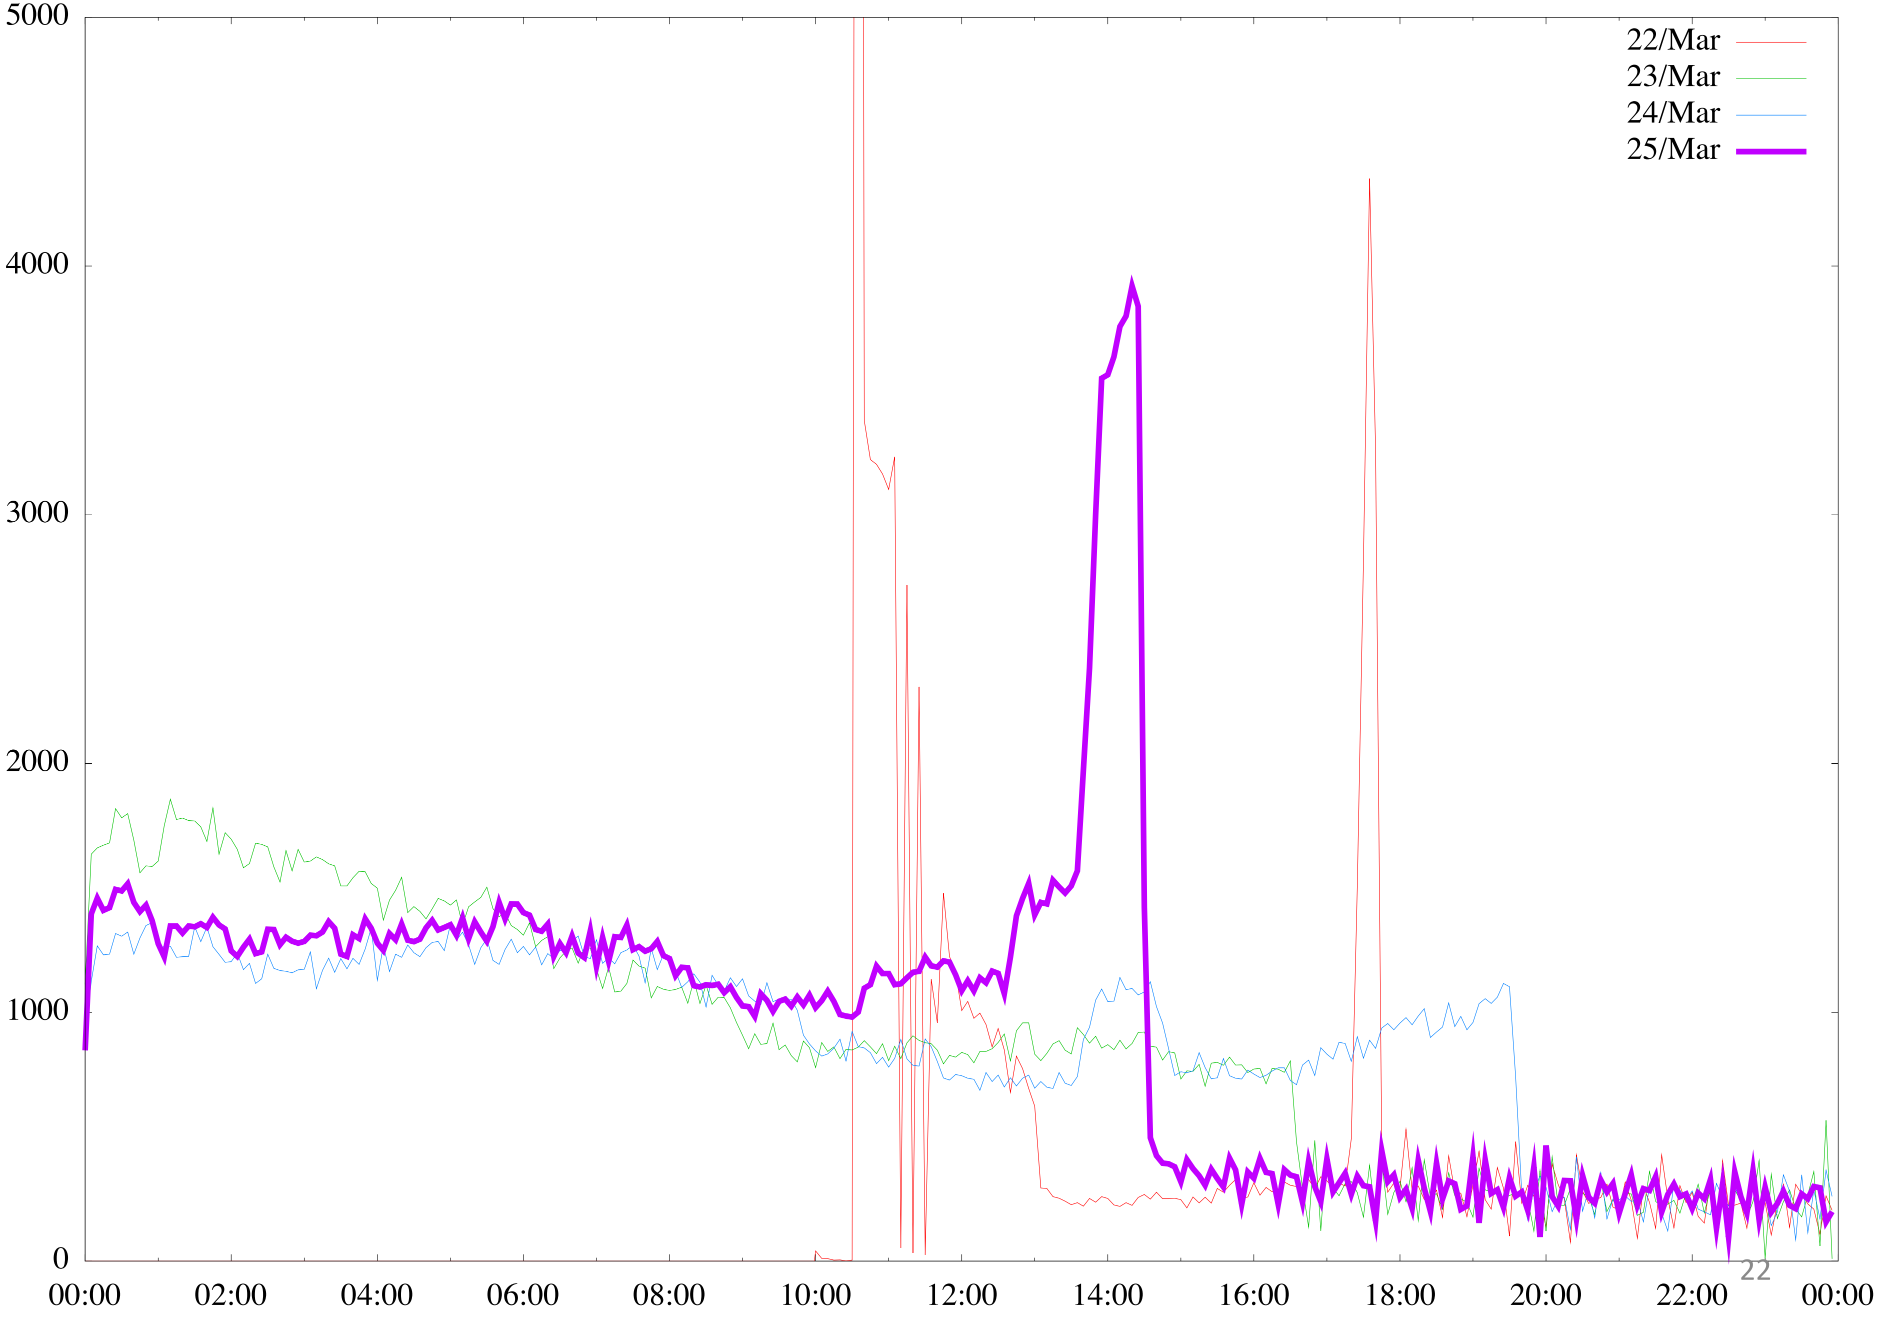

Ad Placement Training – Days 5, 6 & 7

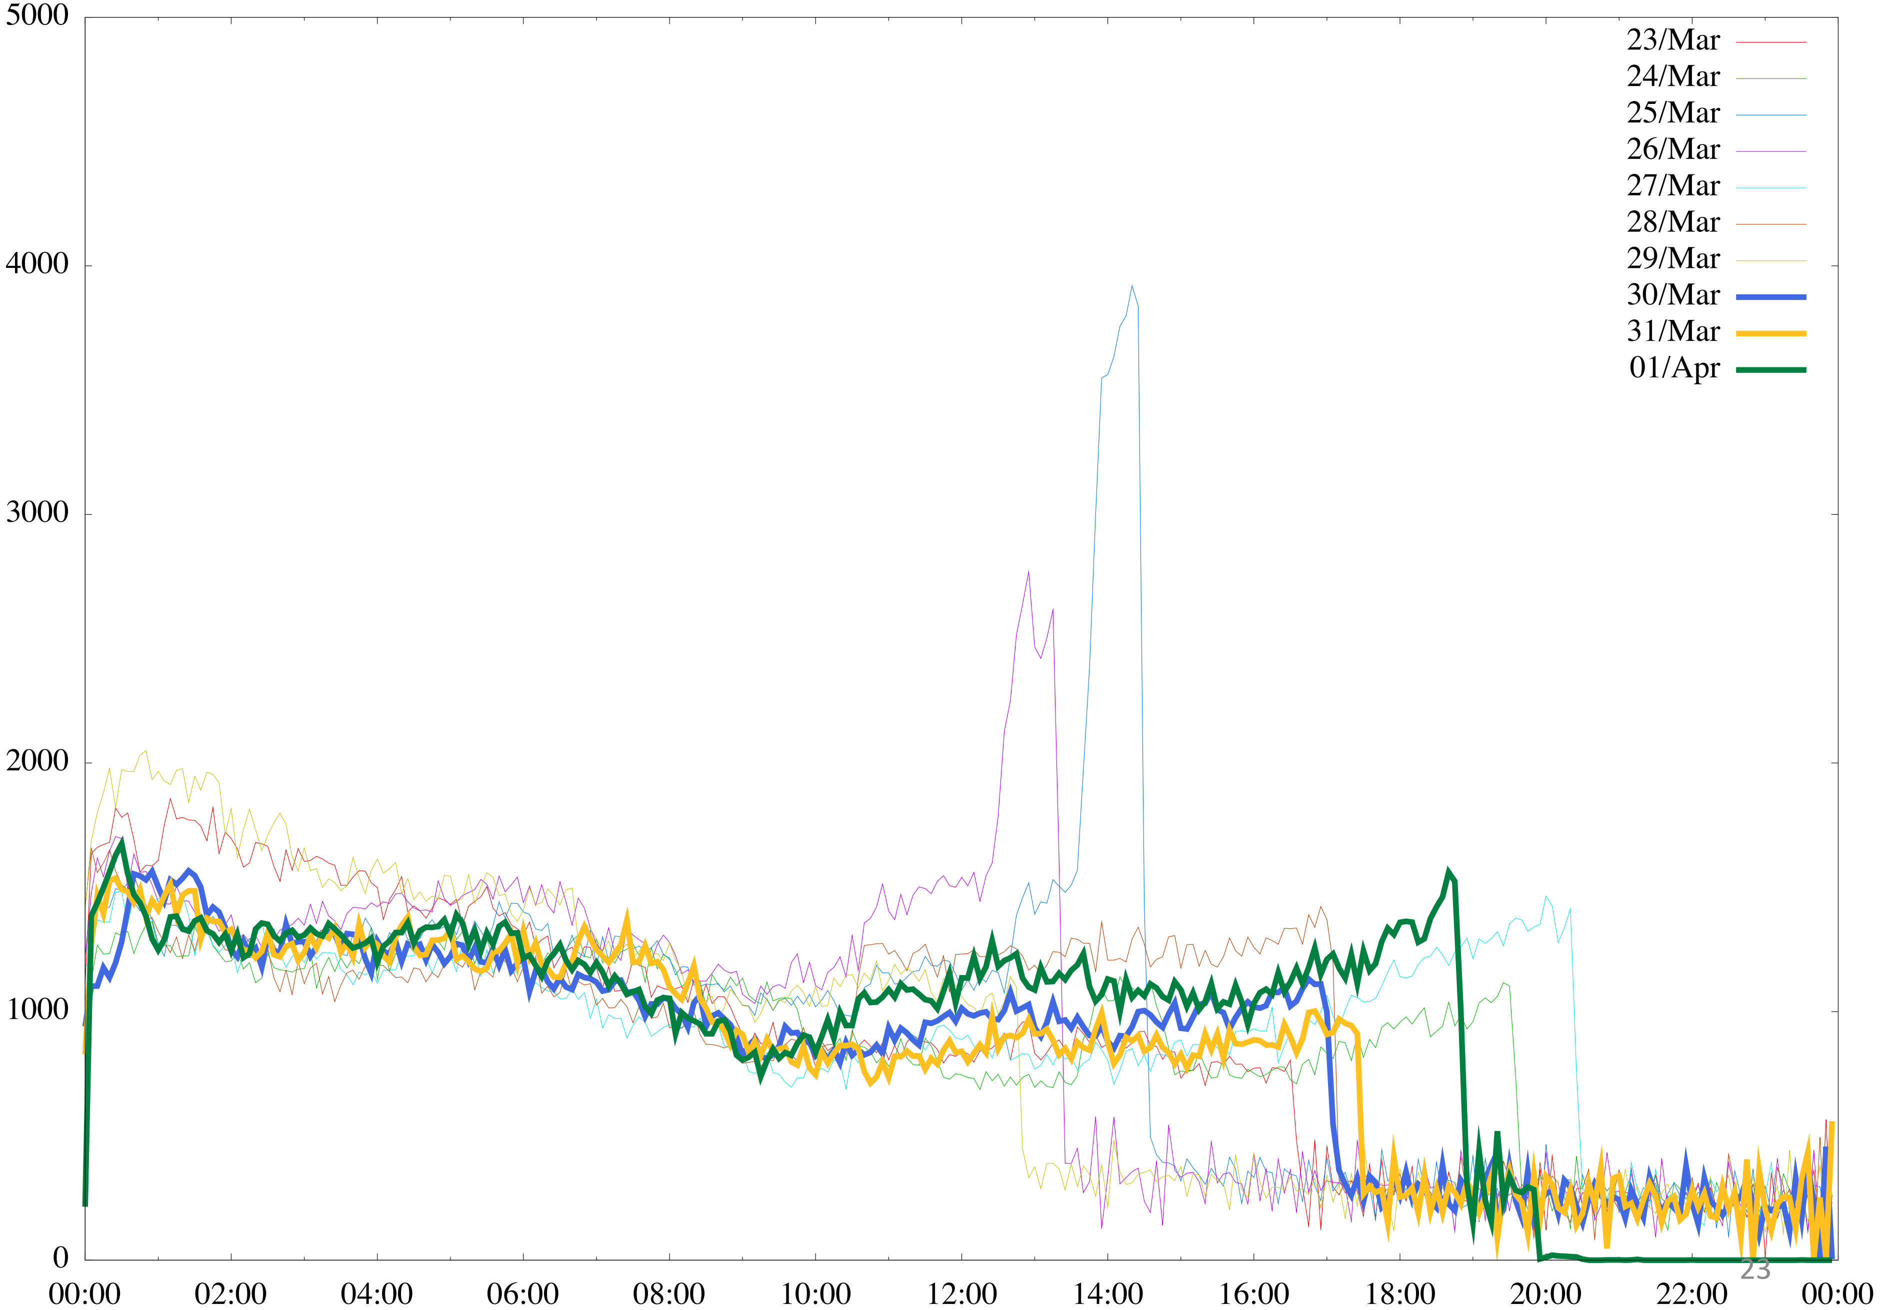

Fresh Eyeballs

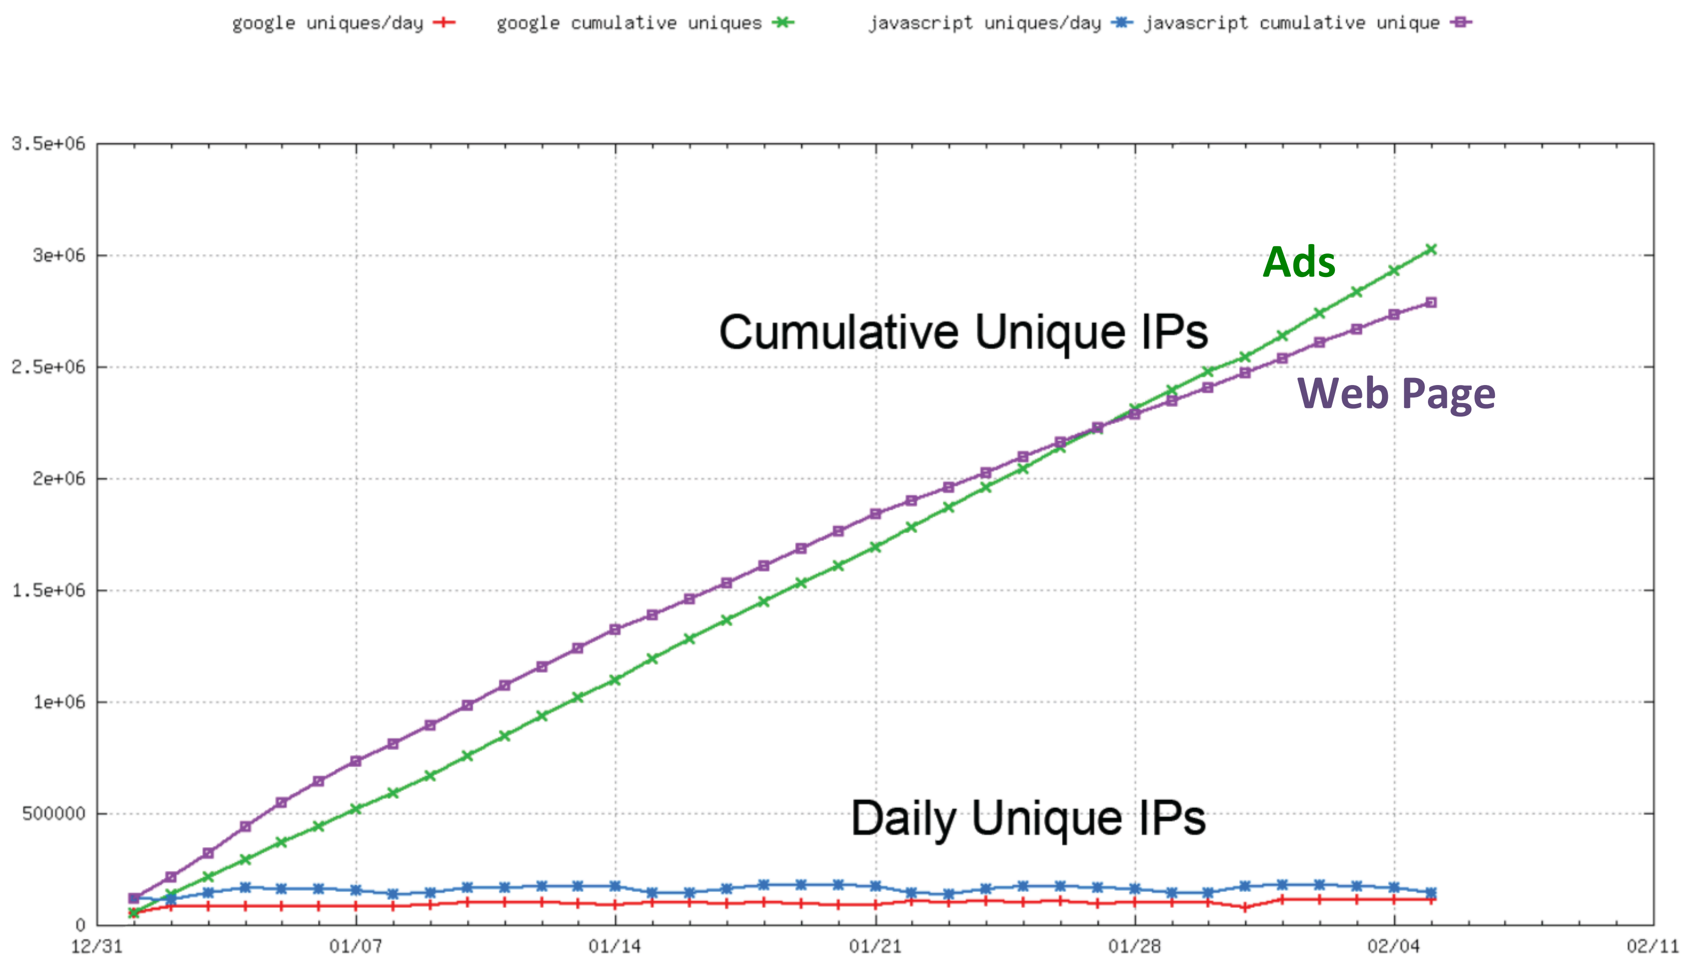

Success!

- 600K – 1M 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 24th May 2015

ITU-T Internet User Census

155,430	VN Vietnam	668,493,485	China
103,517	CN China	282,384872	United States of America
92,107	MX Mexico	252,482905	India
79,092	TH Thailand	110,345878	Brazil
73,702	IN India	109,390190	Japan
65,402	PK Pakistan	87,305661	Russian Federation
64,121	BR Brazil	72,663301	Nigeria
54,637	TR Turkey	71,823404	Indonesia
52,532	US United States of America	71,174958	Germany
52,240	AR Argentina	61,579582	Mexico
48,315	CO Colombia	57,306333	United Kingdom of Great Britain and Northern Ireland
45,216	ID Indonesia	54,114094	France
39,839	PE Peru	45,416941	Iran (Islamic Republic of)
36,962	RU Russian Federation	45,019465	Egypt
34,529	PH Philippines	42,187842	Republic of Korea
33,899	EG Egypt	41,780667	Philippines
22,983	TW Taiwan	40,980368	Vietnam
22,712	RO Romania	39,256999	Bangladesh
22,490	UA Ukraine	35,793673	Italy
22,403	ES Spain	35,503461	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

Weighting the Results

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?

The Generic Approach

- Seed a user's browser with a set of tasks that cause identifiable traffic at instrumented servers
- Rely on unique dns names to ensure that DNS/ Web caching is not used
- The servers collect DNS and Web activity traces that match the URLs in the provided tasks
- Analysis of server logs provides measurement data

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,…

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)

IPv6 Deployment

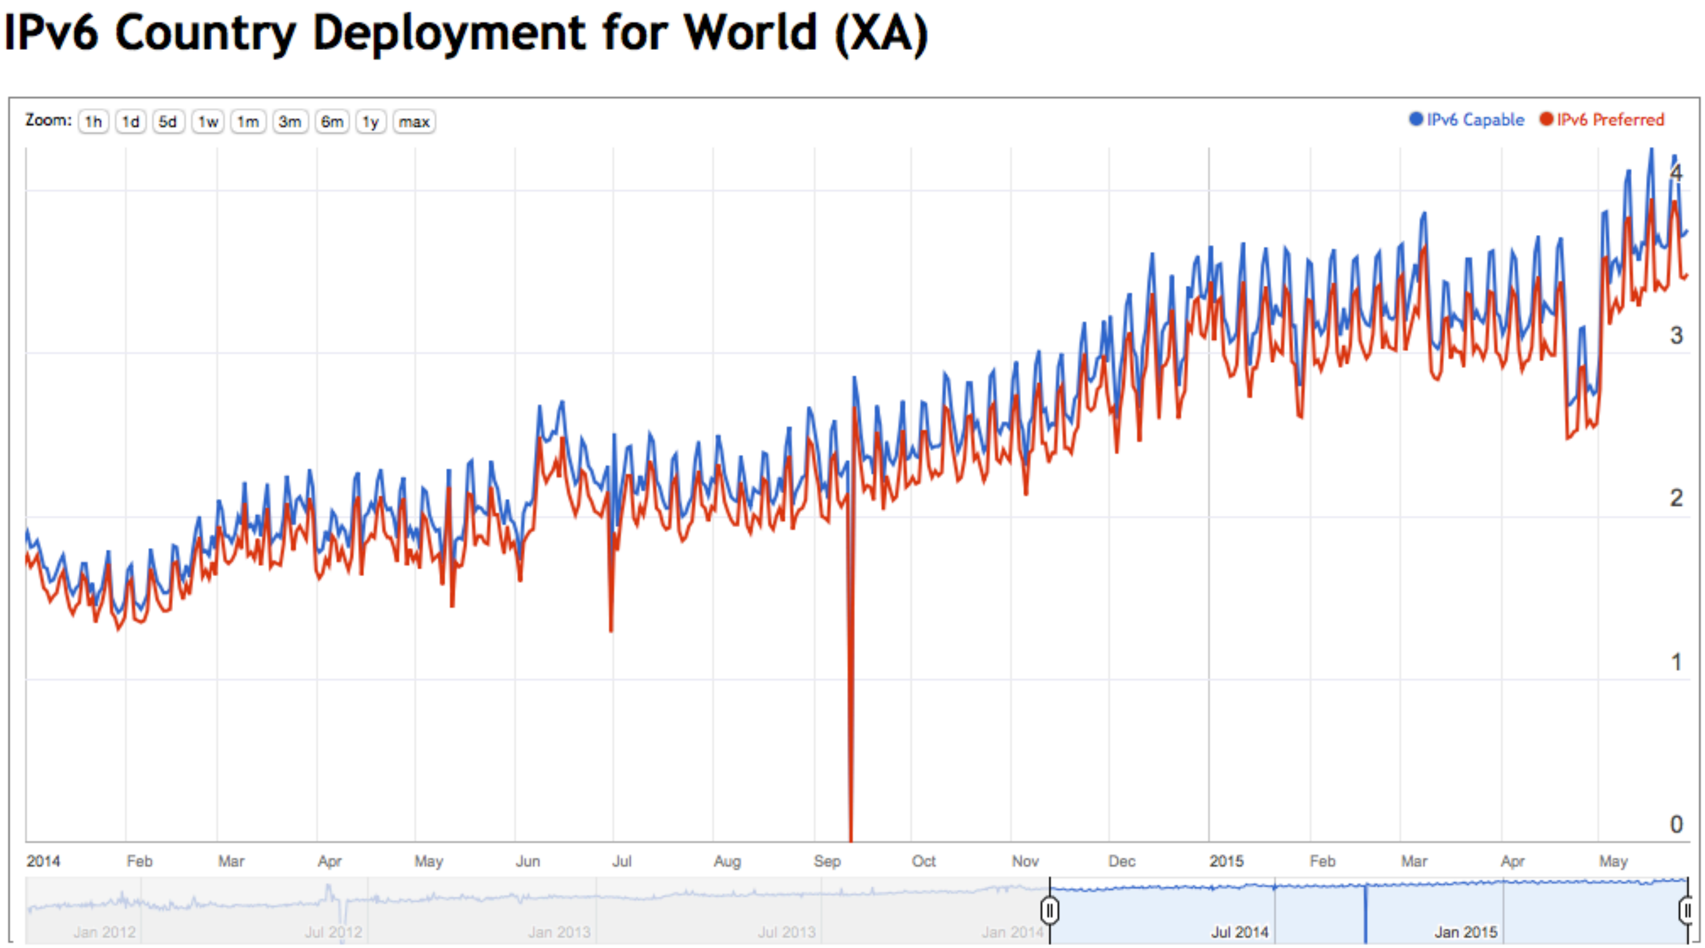

IPv6 Deployment in the US

IPv6 Country Deployment for United States of America (US)

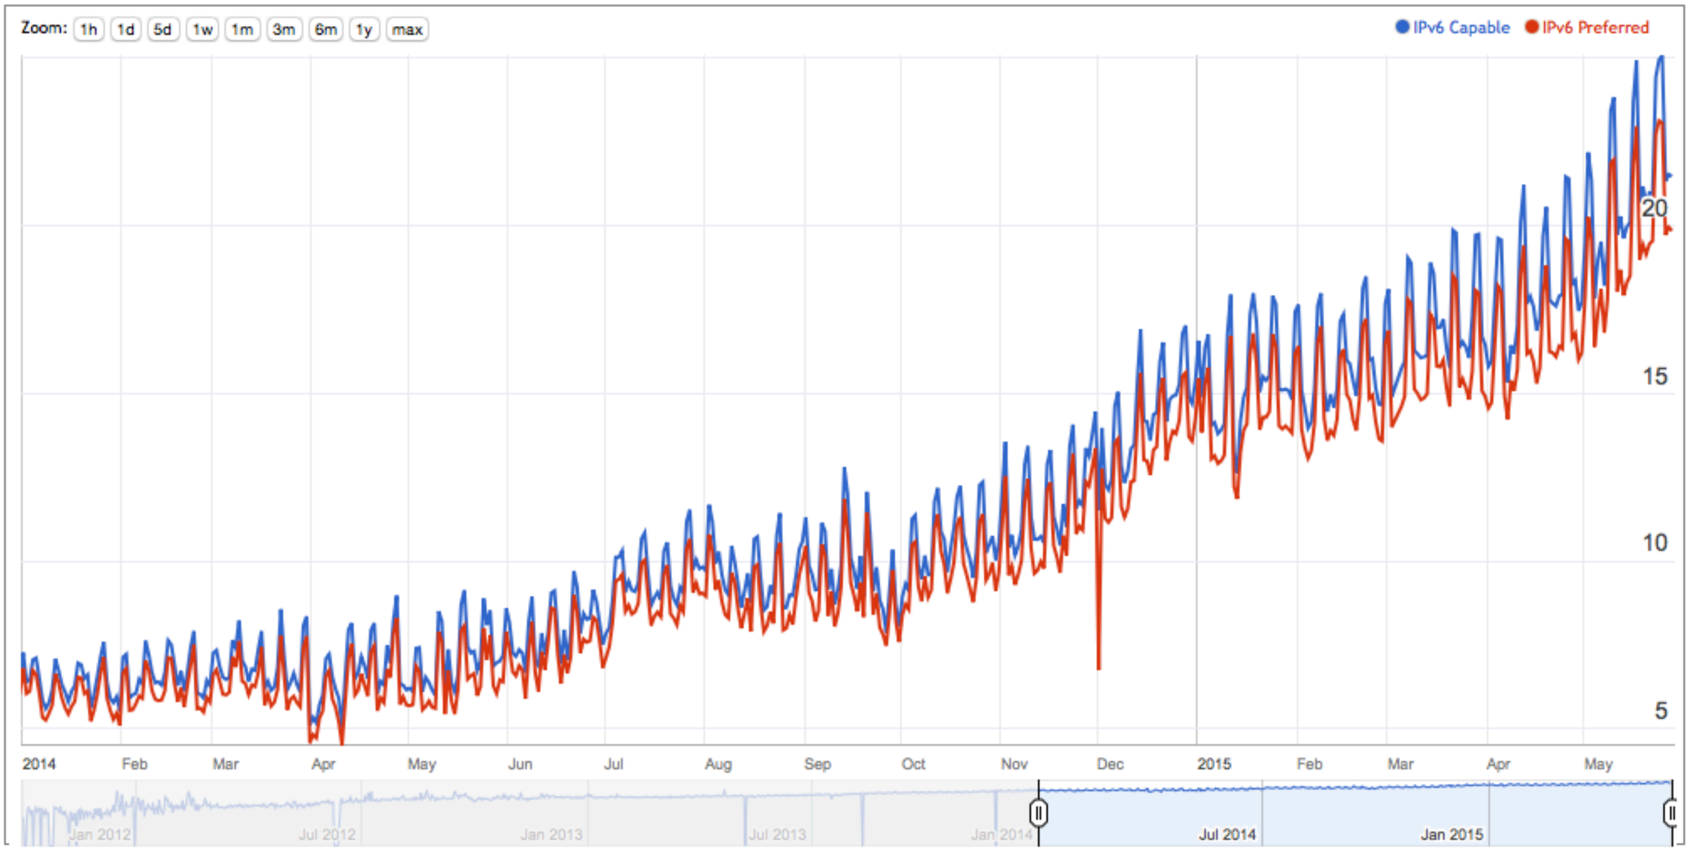

IPv6 Deployment in Comcast

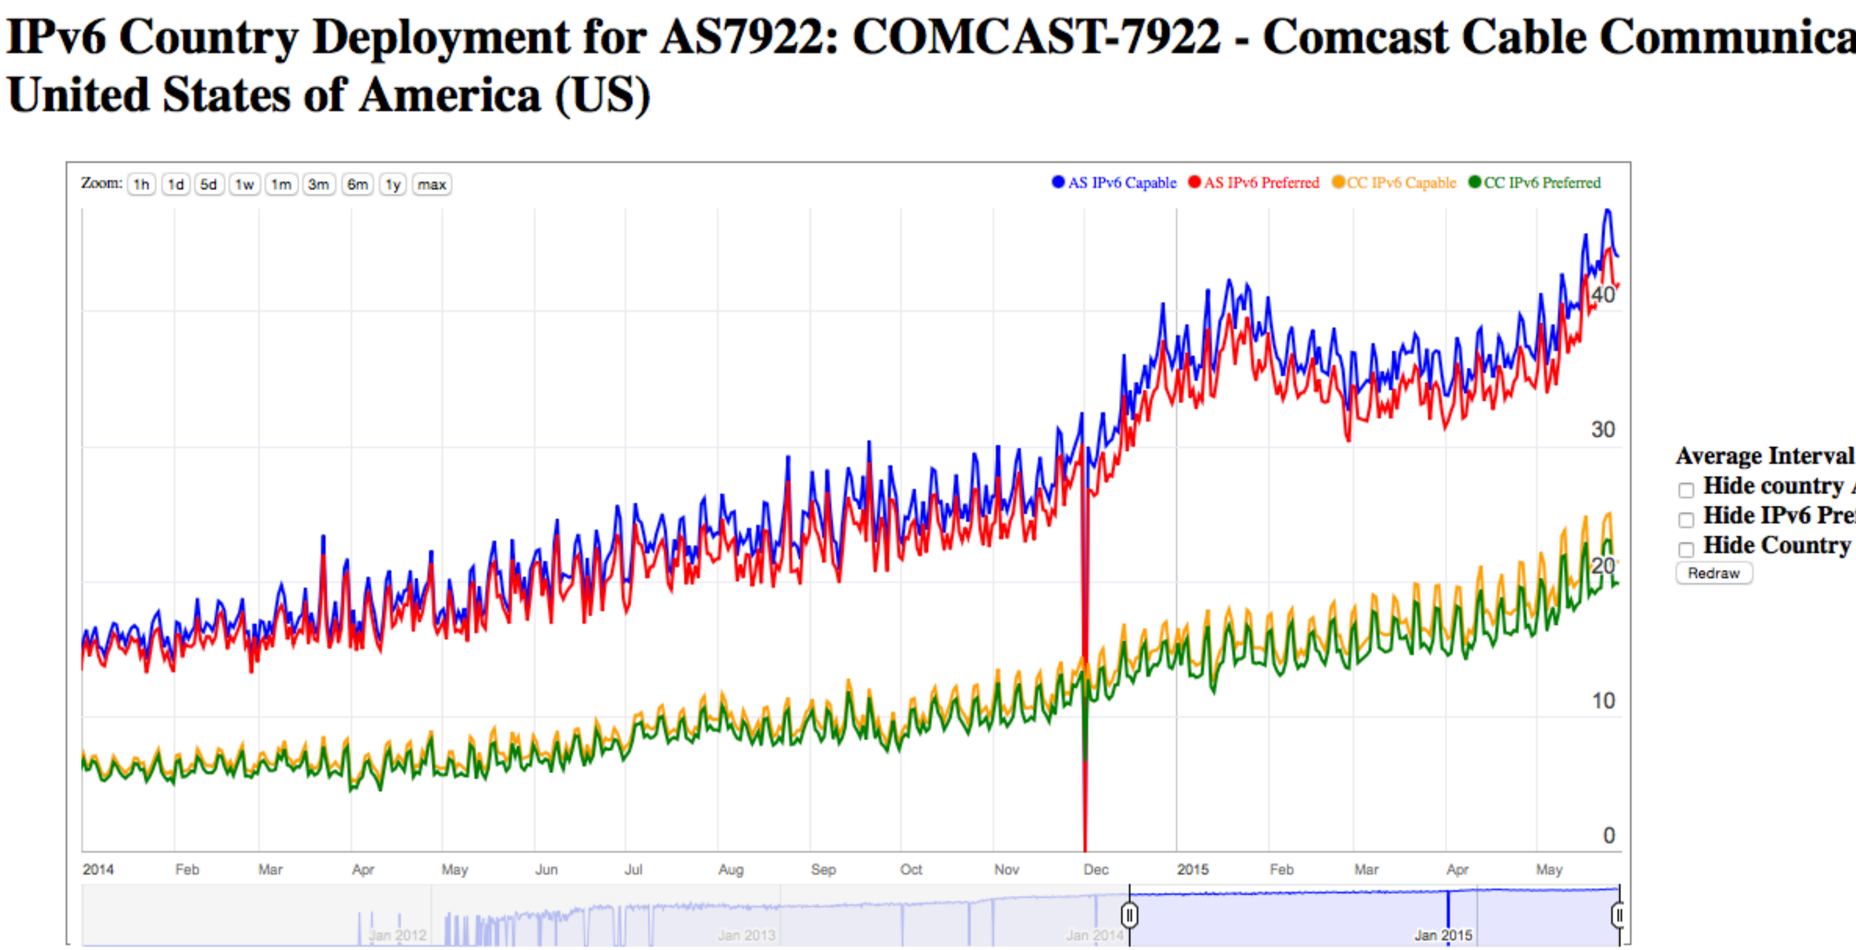

Measuring DNS Behaviours

Understanding DNS behaviour is "messy"

What we would like to think happens in DNS resolution!

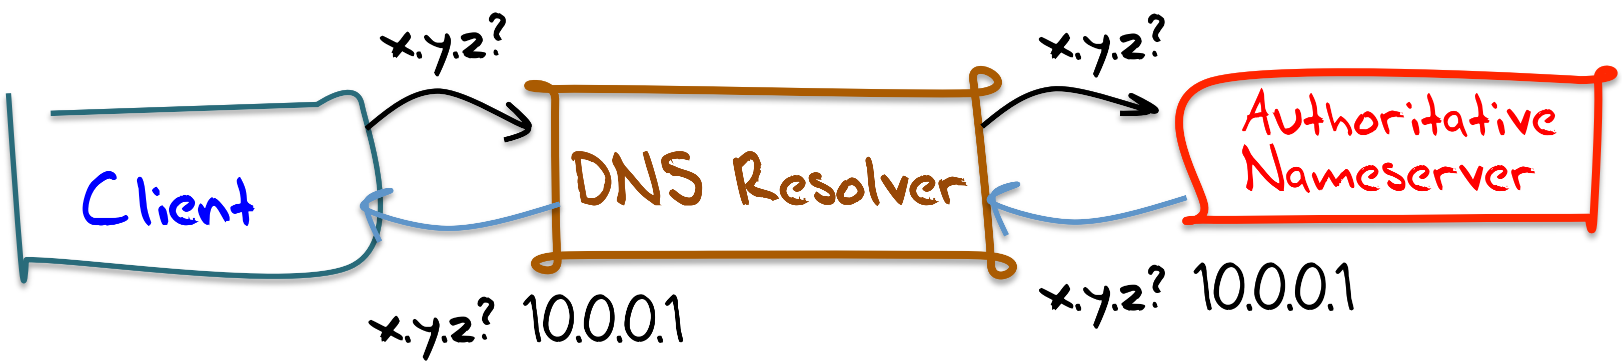

Understanding DNS behaviour is "messy"

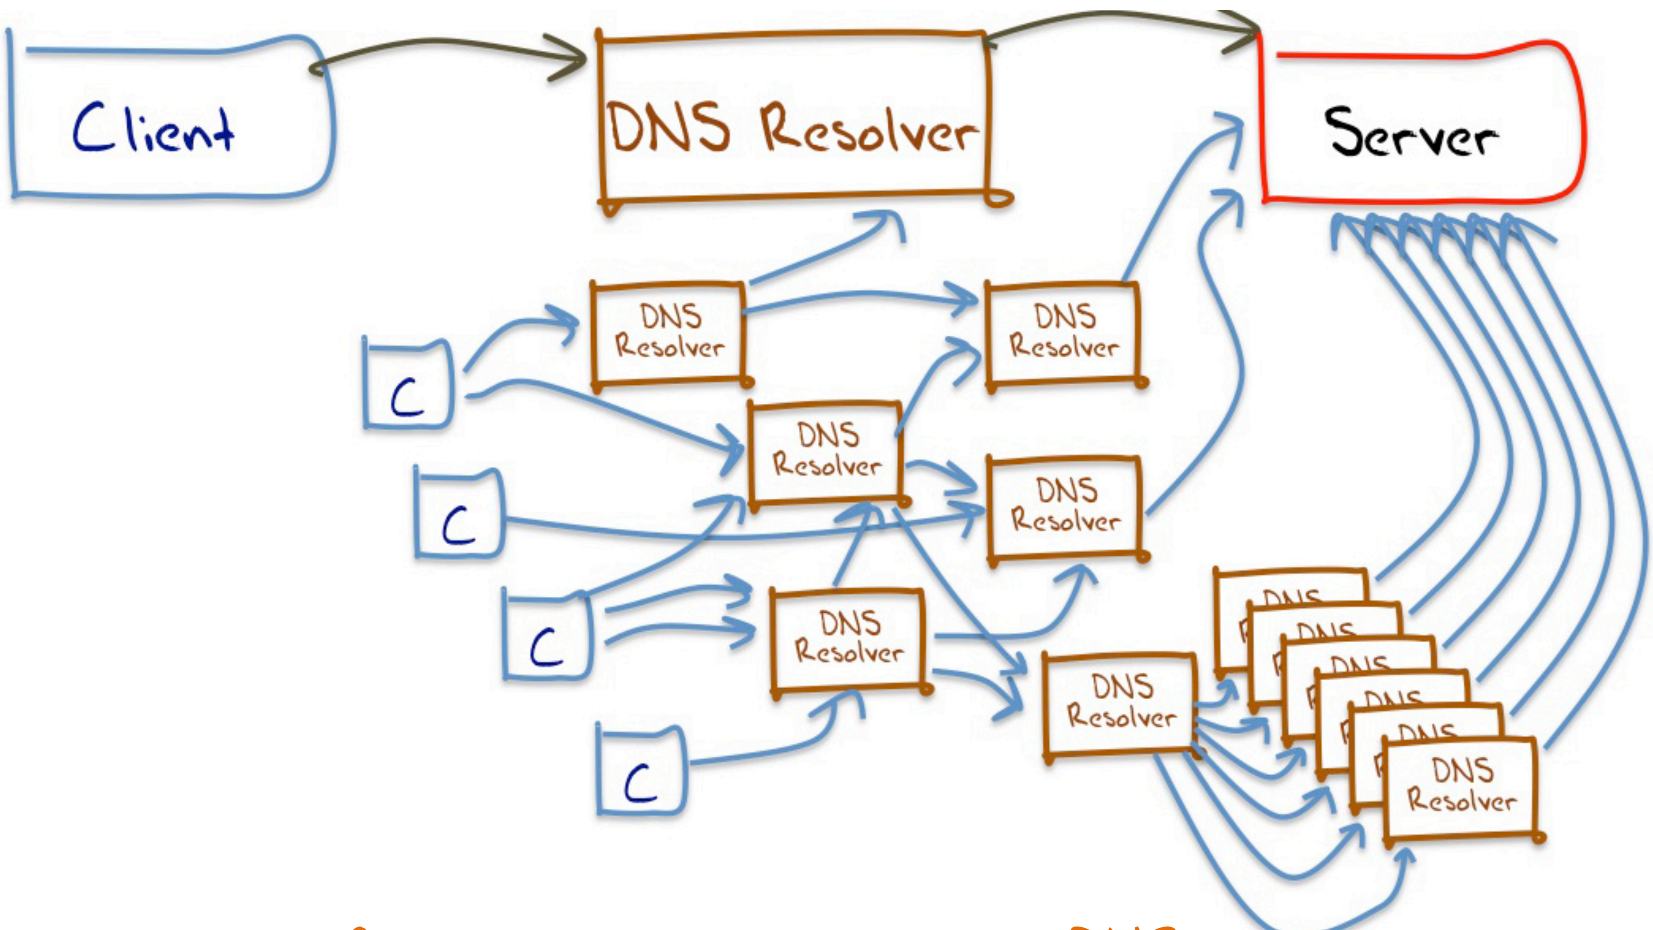

A small sample of what appears to happen in DNS resolution

Understanding DNS behaviour is "messy"

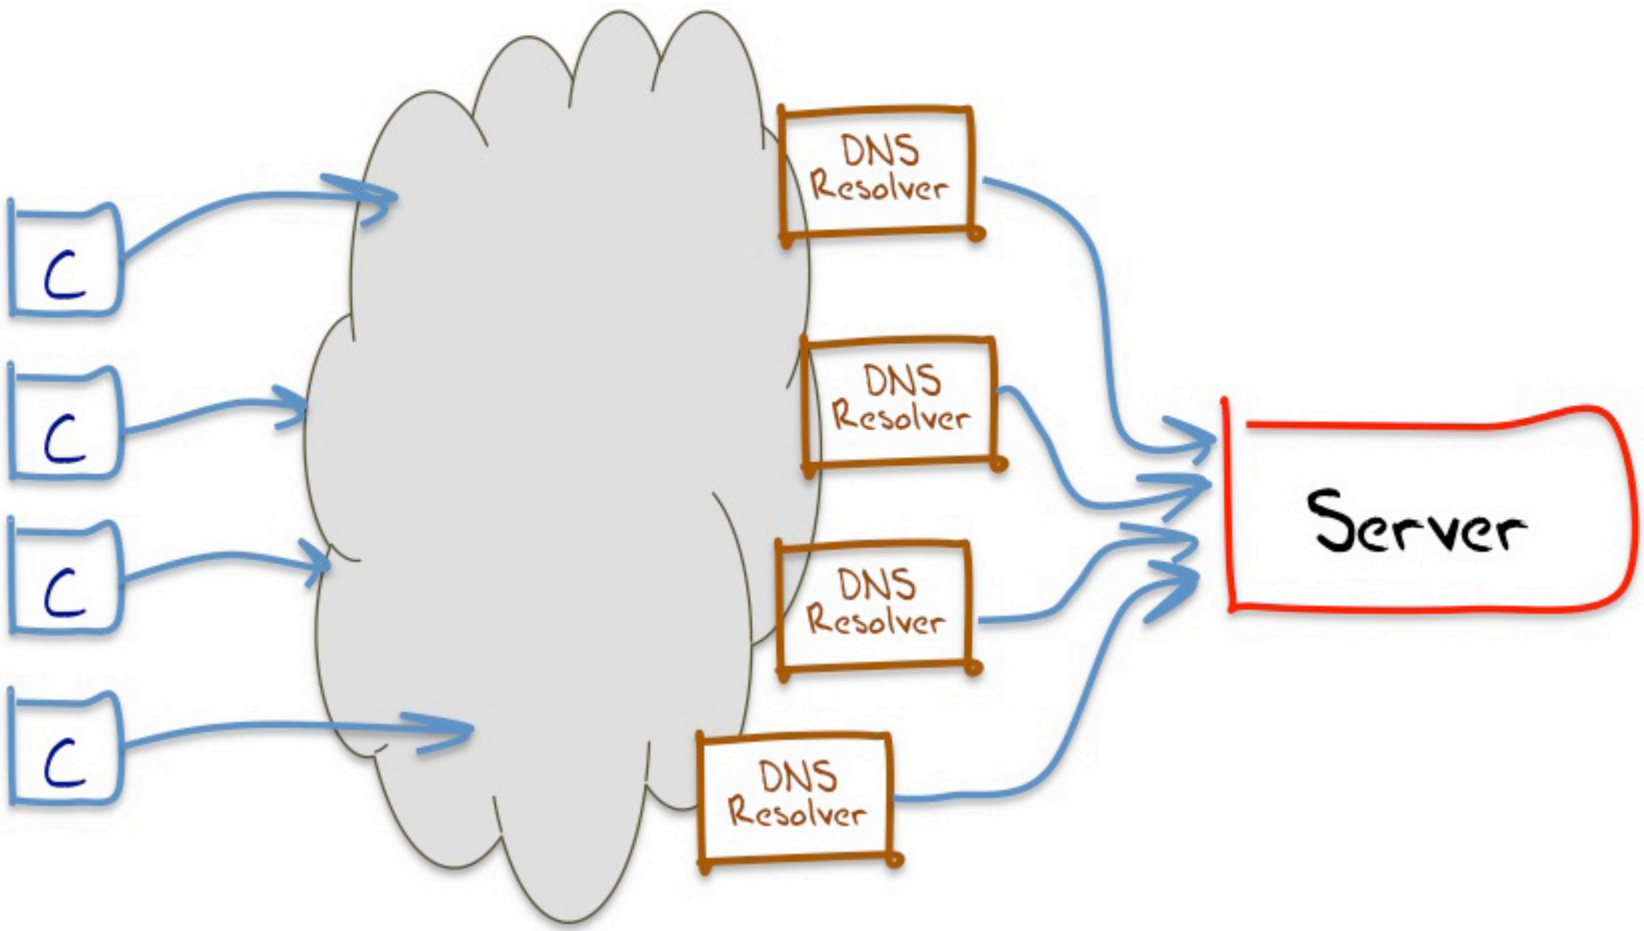

The best model we can use for DNS resolution

This means…

That it is hard to talk about “all resolvers”

- We don’t know the ratio of the number of resolvers we cannot see compared to the resolvers we can see from the perspective of an authoritative name server
- We can only talk about “visible resolvers”

This means…

And there is an added issue with DNSSEC:

– It can be hard to tell the difference between a visible resolver performing DNSSEC validation and an occluded validating resolver performing validation via a visible non-validating forwarder

(Yes, i know it's a subtle distinction, but it makes looking at RESOLVERS difficult!)

This means…

It's easier to talk about **end clients** rather than **resolvers**, and whether these end clients use / don't use a DNS resolution service that performs DNSSEC validation

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

DNSSEC Validation

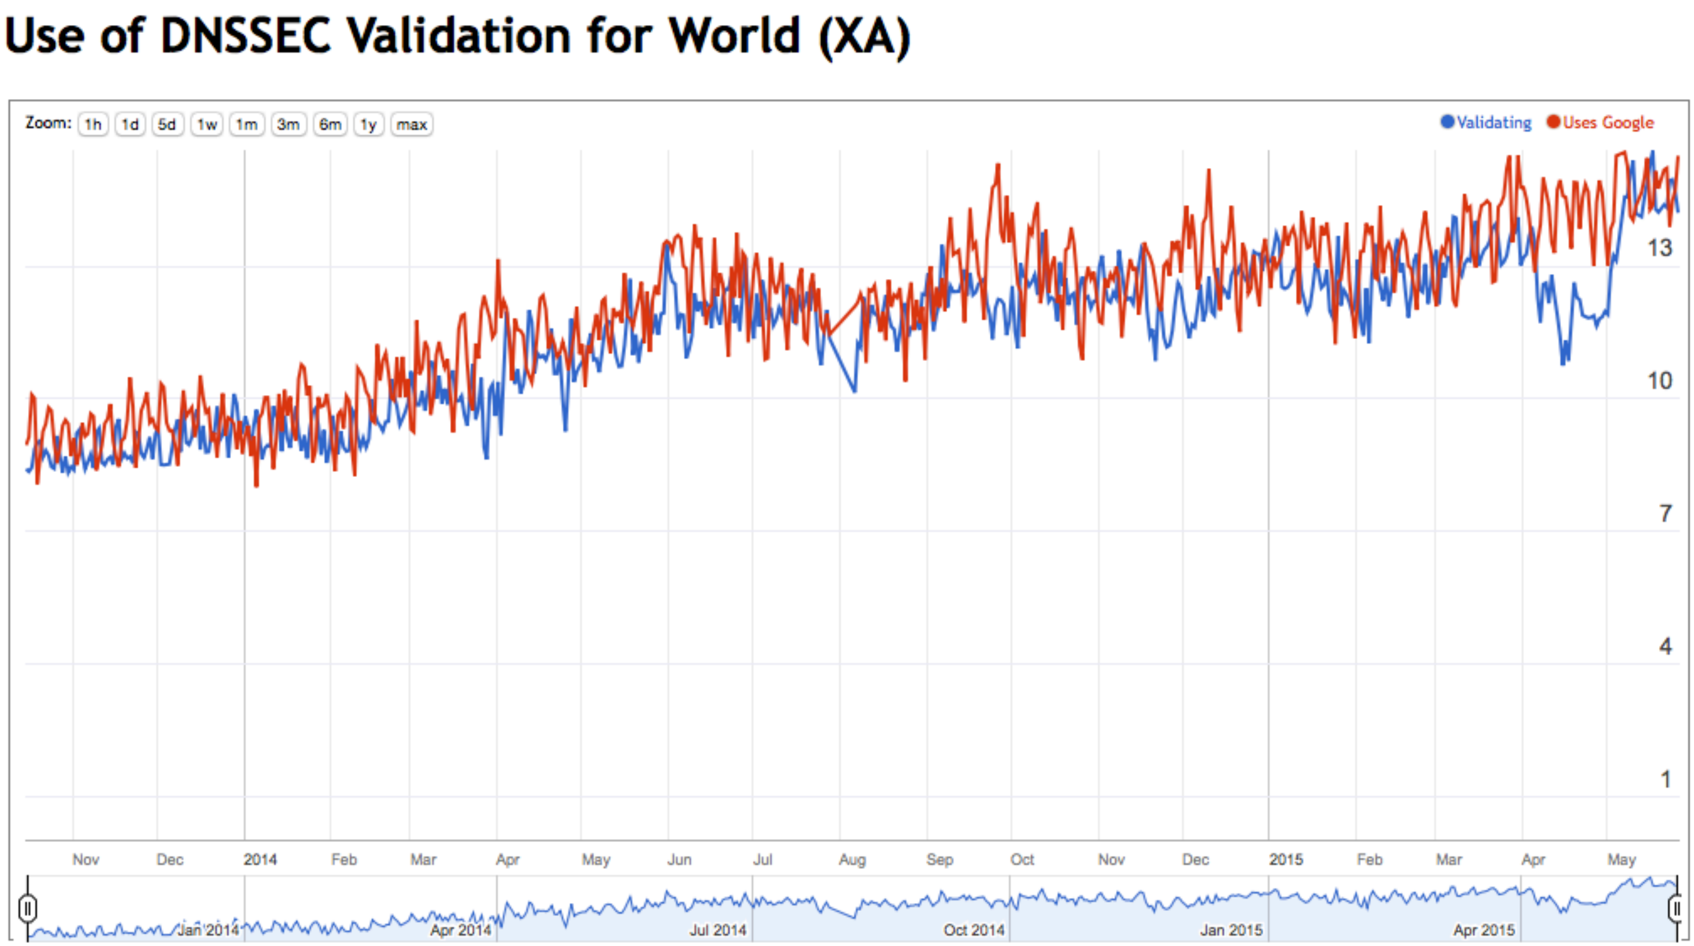

DNSSEC Validation in Sweden

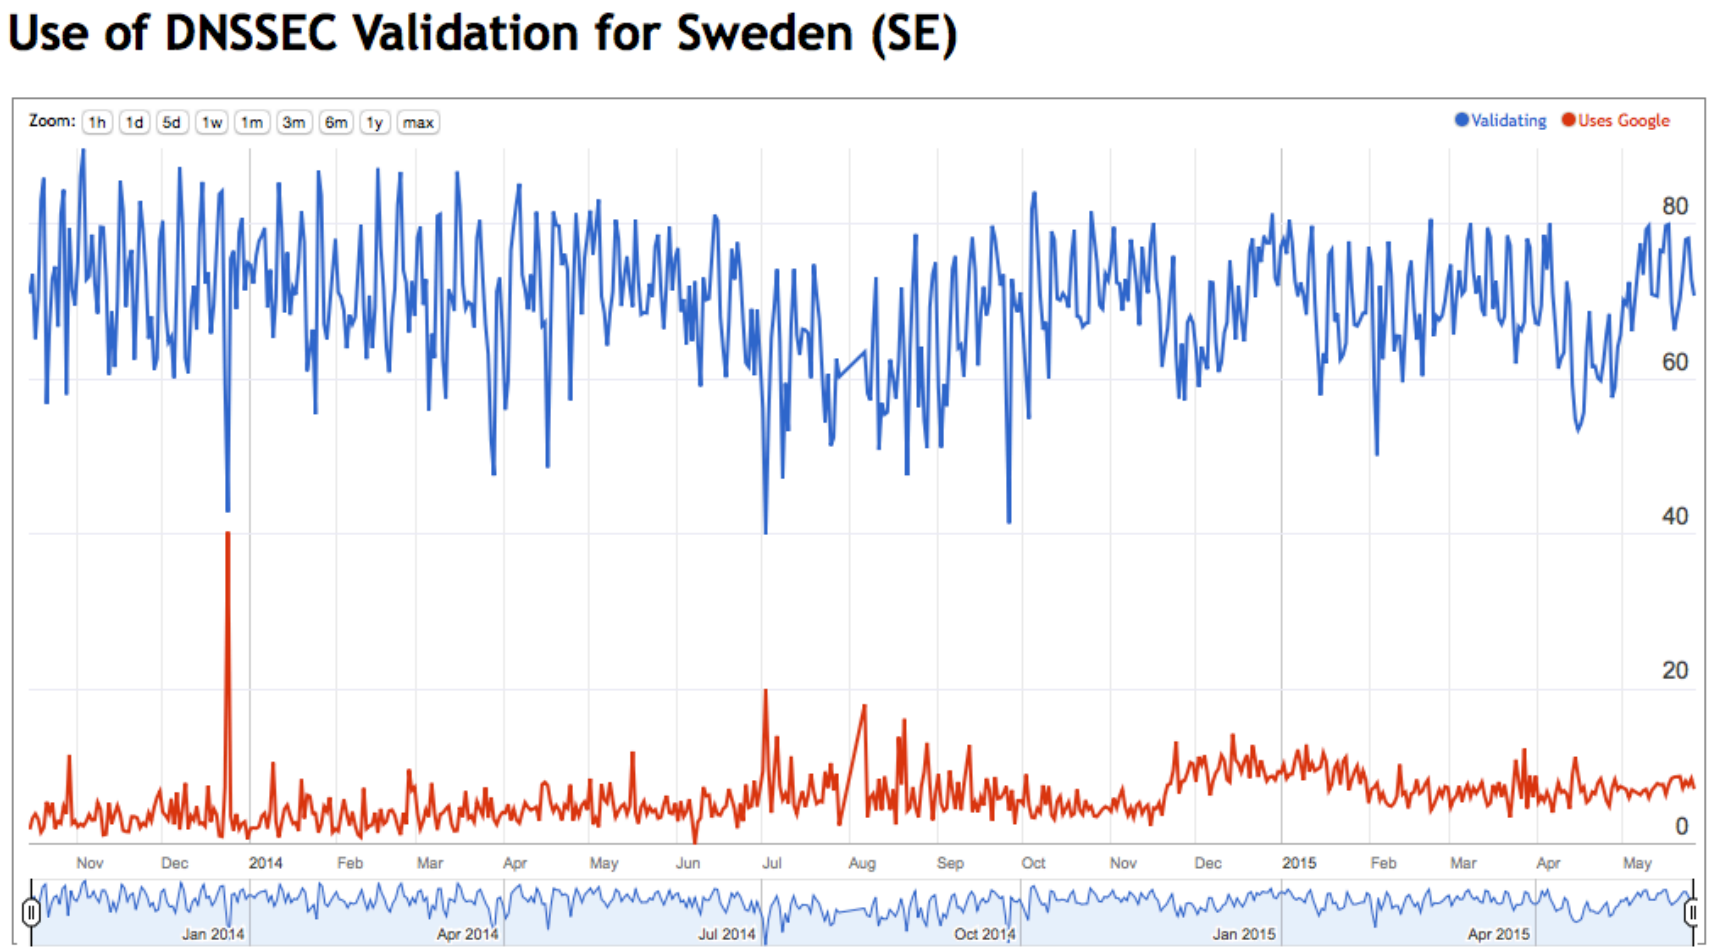

What Else?

- We can isolate the behaviour of individual DNS resolvers using indirection (glueless delegation) within the delegation path
  - How many resolvers fail to resolve a name when the DNS response is 1,444 octets?
  - How may resolvers can use IPv6? How many resolvers prefer to use IPv6?

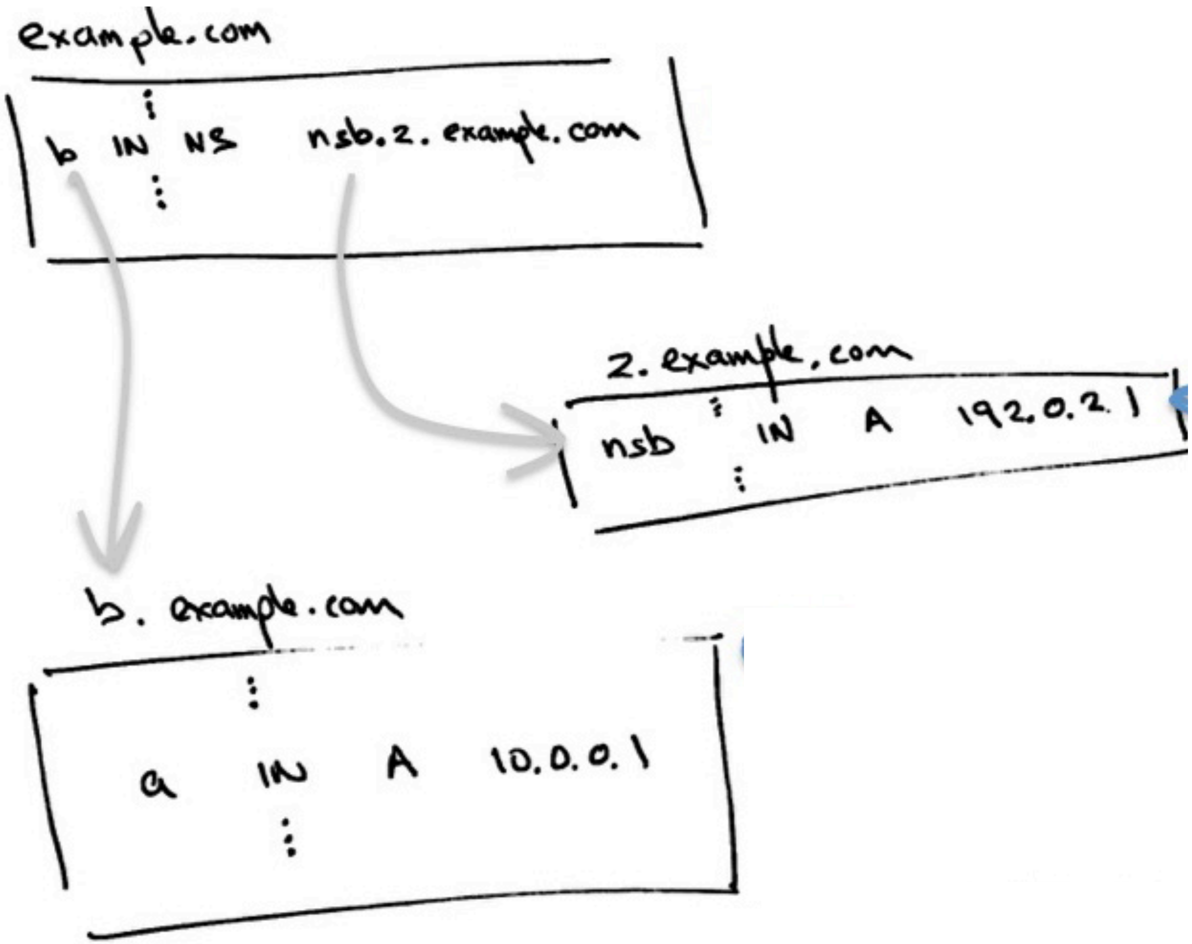

example.com
b IN NS nsb.2.example.com
2.example.com
nsb IN A 192.0.2.1
b.example.com
a IN A 10.0.0.1

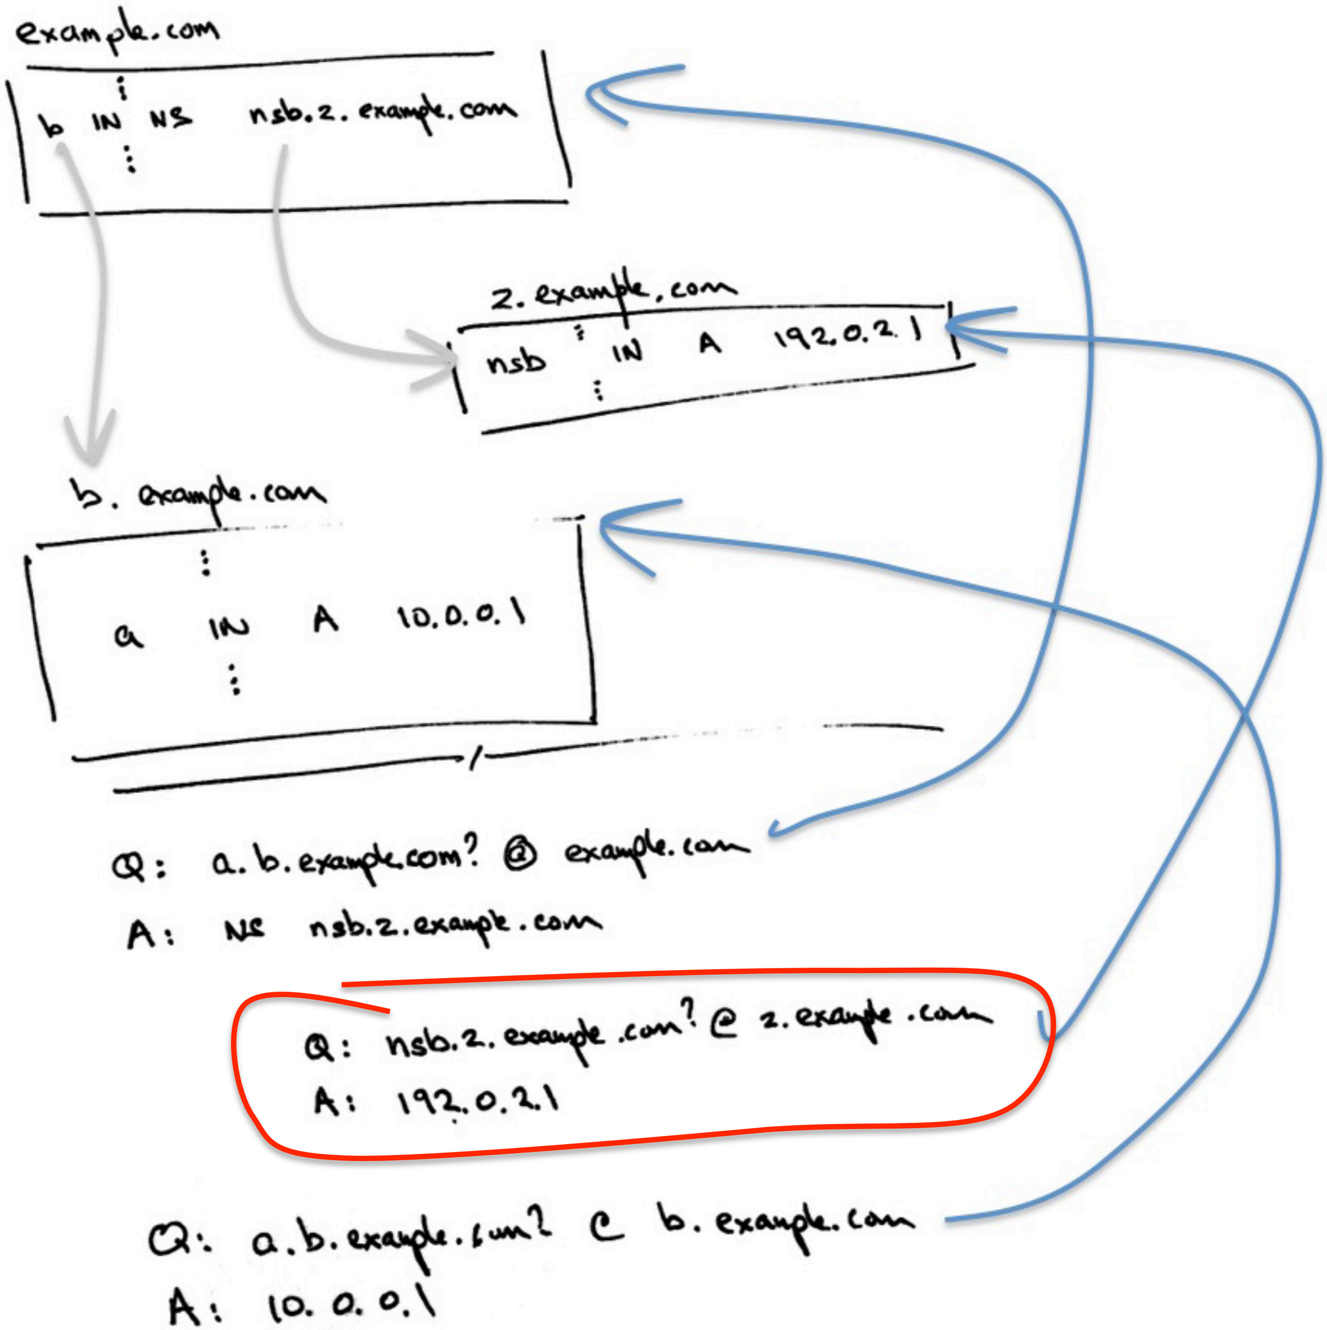

example.com
b IN NS nsb.2.example.com
2.example.com
nsb IN A 192.0.2.1
b.example.com
a IN A 10.0.0.1
Q: a.b.example.com? @ example.com
A: NS nsb.2.example.com
Q: nsb.2.example.com? @ 2.example.com
A: 192.0.2.1
Q: a.b.example.com? @ b.example.com
A: 10.0.0.1

What Else?

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

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?

What Else?

- Digital Stalking: We deliver a unique URL to a single end device via the AD placement mechanism
  - We expected that the script would be executed once.
  - But for some 2% of users we see the script executed a second time!

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…

- Its 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

Where now?

- We need to move this entire test system to use TLS
  - Too much malware is trying to intrude on the ad delivery system (i.e. the Great Canon!)
  - Ad delivery systems are pushing to secure any third party references
- We need to migrate the entire scripting system from Flash to an HTML5 base
- We need to migrate to use a customized DNS server that performs a combination of pseudo zone creation and on-the-fly signing
- We are moving off Apache to NGINX
- We need to improve our server infrastructure in location and capacity

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

Questions?

APNIC Labs:

Geoff Huston research@apnic.net