

Measuring DNS over IPv6

Geoff Huston AM
APNIC Labs

RFC Recommendations

RFC3901 – September 2004 “DNS IPv6 Transport Operational Guidelines”:

- Every recursive name server SHOULD be either IPv4-only or dual stack
- Every DNS zone SHOULD be served by at least one IPv4-reachable name server

*Which is saying as an IPv6 Operational guideline “you better keep IPv4 going”
This RFC actually says very little about IPv6!*

Proposed: 3901bis

Current IETF draft proposed to update RFC3901 by saying:

- It is RECOMMENDED that at least two NS for a zone are dual stack name servers
- Every authoritative DNS zone SHOULD be served by at least one IPv6 reachable authoritative name server

Which is saying as an IPv6 Operational guideline "time to take IPv6 seriously" and NOT saying that servers need to keep IPv4 around— which is largely the opposite of the advice in RFC 3901!

The assumption behind 3901bis

That IPv6 is now a mature and well understood technology,
and using IPv6 as the transport for the DNS is as efficient and
as fast as using IPv4

The assumption behind 3901bis

That IPv6 is now a mature and well understood technology,
and using IPv6 as the transport for the DNS is as efficient and
as fast as IPv4

Is this true in today's Internet?

IPv6 and the DNS

How widely is the use of IPv6 supported in the Internet's DNS?

- If you placed authoritative servers on an IPv6-only service how many users would be able to reach you?

Measuring the DNS

- Is not as straightforward as it might sound...

Measuring DNS Resolver Behaviour

We really don't understand what a "resolver" is!

- It could be a single platform running an instance of DNS resolver code
- It could be a collection of independent back-end systems with a load distributor front end facing clients
- It could be a hybrid collection where the back ends synchronise each other to emulate a common cache
- It is a stub, recursive, or forwarding resolver, or a hybrid behaviour distributed across multiple resolution systems
- A resolver may have 1 client, or millions of clients, or anything in between

When we talk about **DNS Resolvers** it's challenging to understand exactly what we are talking about!

Measuring "DNS Queries"

We don't understand what "a query" is!

Which sounds silly, but the distributed resolution process causes a 'fan out' of queries as part of the resolution process when a single query may cause a number of 'discovery' queries to establish the identity of the authoritative server(s) for the name


Resolvers all use their own timers for retransmission

Queries have no "hop count" or "resolver path" attached - there is no context to understand the reason for a query

Queries appear to have a life of their own!

Also, the DNS is noisy!


The DNS cannot use multiple query types in the one DNS transaction – if an application wants to establish the IPv4 address, the IPv6 address, and the application level protocol for service (HTTPS query), then that's three distinct queries in today's DNS



The DNS is VERY noisy!


There are many repeated DNS queries:

- 25% - 35% of query name + type are repeated 2 or more times for the A / AAAA / HTTPS query types
- The use of UDP transport motivates approaches of rapid-fire query fan-out to compensate for unpredictable response times



APNIC's DNS Experimental Rig

- We use Google's ad network to "seed" DNS queries
- We make parts of the DNS name unique to each measurement
 - That way the recursive resolvers have no cached data and are forced to query the authoritative server
- We observe the recursive-to-authoritative query process by instrumenting the authoritative server, and match experiment placement records to the server's DNS logs



Back to the measurement question


IPv6 and the DNS

How widely is the use of IPv6 supported in the Internet's DNS?

- If you placed authoritative servers on an IPv6-only service how many users would be able to reach you?

Use of IPv6 Transport in the DNS

IPv6 DNS Use in World (XA)



How "good" is this measurement?

- We place the 1x1 web blot behind a unique DNS name whose only Authoritative Server can only be queried by using IPv6 transport
 - The 1x1 blot is itself on a dual stack server and the DNS zone contains both A and AAAA records
 - But this DNS resolution is only accessible if the resolver can pose the DNS query over IPv6
- So, it looks like a reasonable measurement – right?

But

- We are counting the **absence** of data – i.e. the number of sample points who do **not** fetch a 1x1 blot
- Some users may not even attempt to make the DNS query – the user may terminate the ad script early
- Some users may successfully query the DNS over IPv6, yet fail to perform the 1x1 blot fetch
 - Should we count those DNS-query-only users as a positive result?
 - But what if the DNS result is being filtered by filtering middleware?

A DNS-base Measurement

Can we perform this same measurement of IPv6 capability in the DNS using only the DNS?

A DNS-base Measurement

Can we perform this same measurement of IPv6 capability in the DNS using only the DNS?

- Yes, by using a “glueless delegation”

Glueless Delegation in the DNS

Parent Zone: example.com

```
...  
a    NS    nsa.sibling.zone.  
...
```

Query: a.example.com to example.com server
Resp: Referral, name server *nsa.sibling.com*
(with no attached glue for this NS name)

Glueless Delegation in the DNS

Parent Zone: example.com

```
...  
a    NS    nsa.sibling.zone.  
...
```

Sibling Zone: sibling.zone

```
.    NS    ns.sibling.zone.  
ns   AAAA  2001:db8::1  
...  
nsa  AAAA  2001:db8::2  
...
```

IPv6 Only

Query: nsa.sibling.zone to ns.sibling.zone server

Resp: IP AAAA records for *nsa.sibling.com*
(*the name server is IPv6 only*)

Glueless Delegation in the DNS

Parent Zone: example.com

```
...  
a NS nsa.sibling.zone.  
...
```

Sibling Zone: sibling.zone

```
. NS ns.sibling.zone.  
ns AAAA 2001:db8::1  
...  
nsa AAAA 2001:db8::2  
...
```

IPv6 Only

Child Zone: a.example.com

```
. NS nsa.sibling.zone.  
...  
. A 192.0.2.2  
  AAAA 2001:db8::2
```

IPv6-only

Query: a.example.com to nsa.sibling.zone server
Resp: IP A/AAAA records for *a.example.com*

Glueless Delegation in the DNS

Parent Zone: example.com

```
...  
a NS nsa.sibling.zone.  
...
```

Sibling Zone: sibling.zone

```
. NS ns.sibling.zone.  
ns AAAA 2001:db8::1  
...  
nsa AAAA 2001:db8::2
```

IPv6 Only

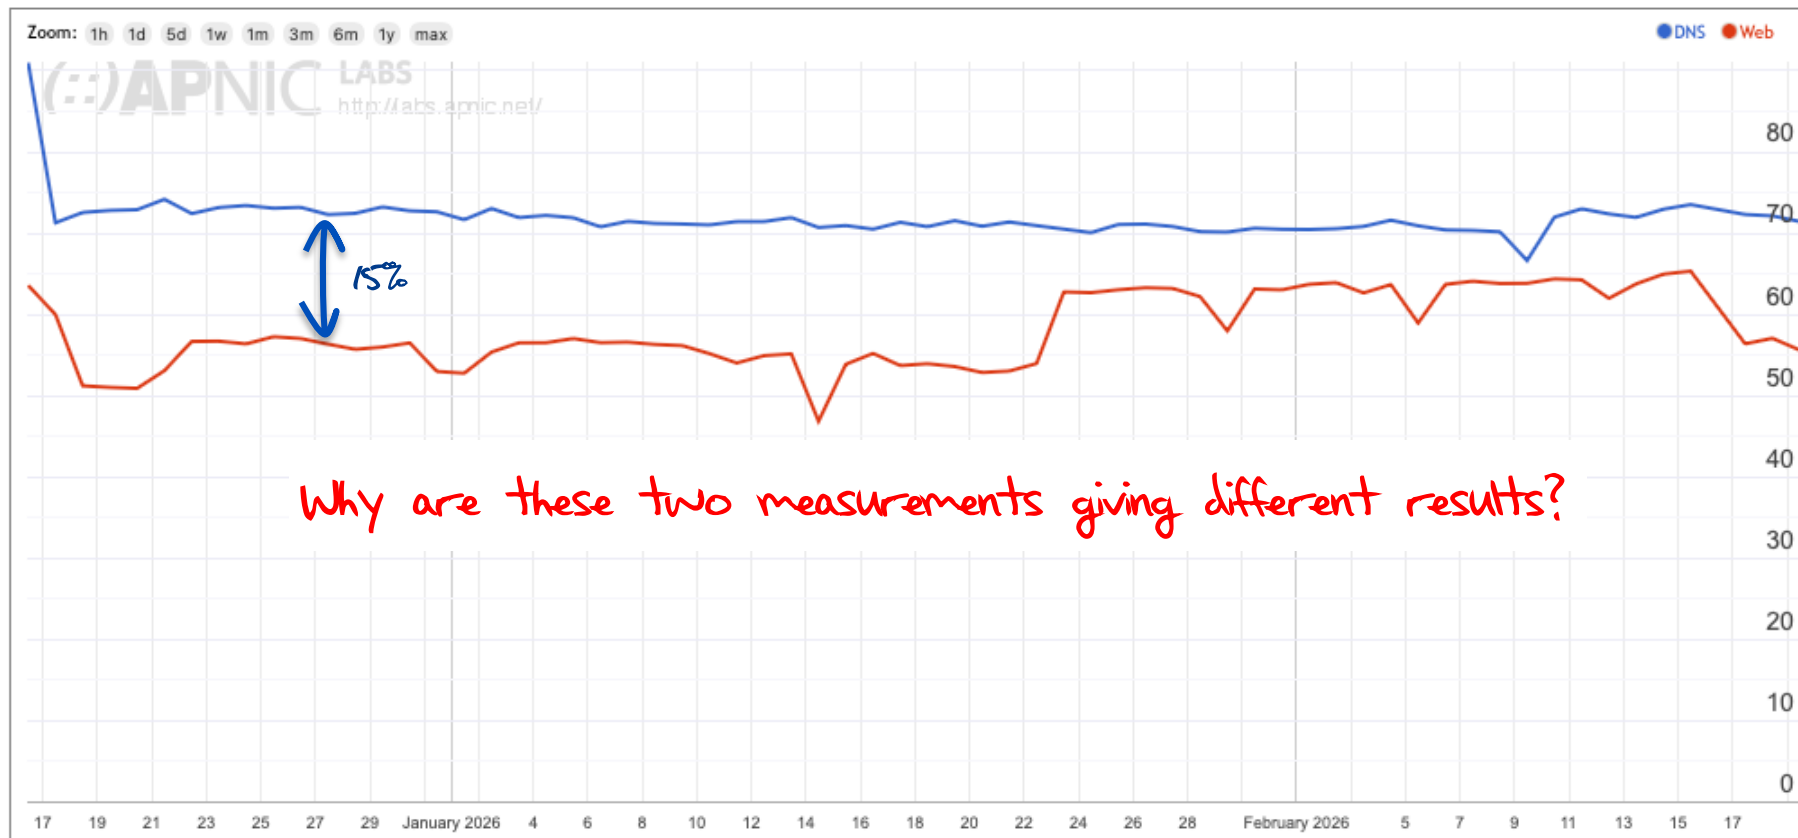
If you can query the child zone, then your DNS resolver environment can make queries over IPv6 and can receive responses

Child Zone: a.example.com

```
. NS nsa.sibling.zone.  
...  
. A 192.0.2.2  
AAAA 2001:db8::2
```

IPv6-only

Comparing the two measurements

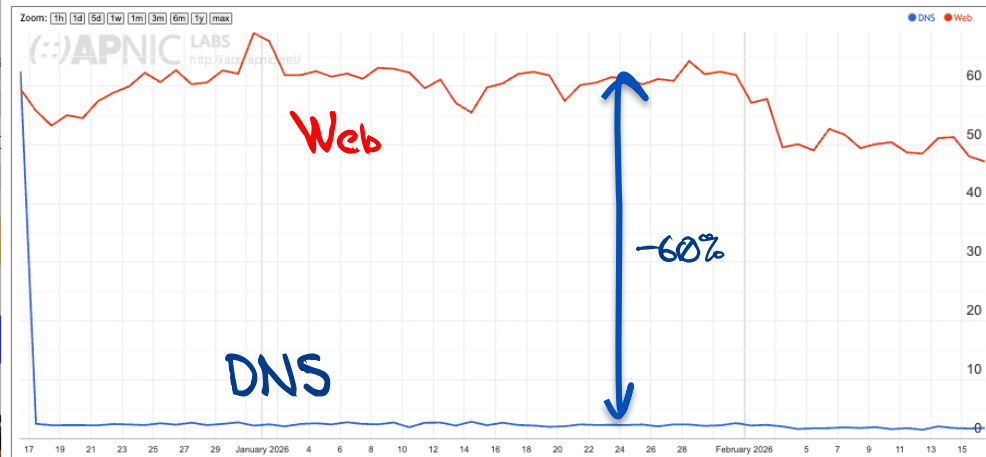


Why are they different?


- There is an expectation that the glueless resolution of a name is the DNS does not rely on application behaviour – the entire measurement is undertaken within the DNS's recursive resolution process, so the client's stub resolver cannot abort the DNS resolution process.
- We expect that the DNS result would be “higher” due to the lossy nature of the application's DNS to Web fetch transition. The global average of this difference is around 10%.
- But averages can hide anomalies...

Measurement Anomalies!

IPv6 DNS Use in Algeria (DZ)



IPv6 DNS Use in Libya (LY)

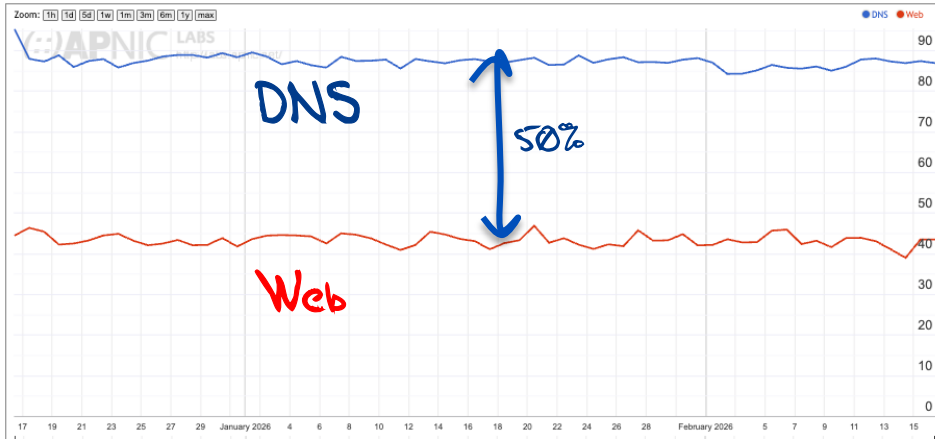


IPv6 DNS Use in Egypt (EG)

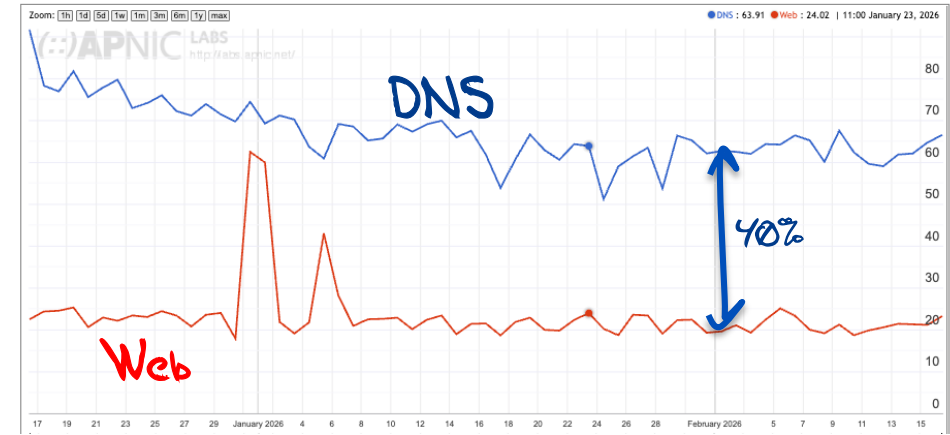


More Anomalies

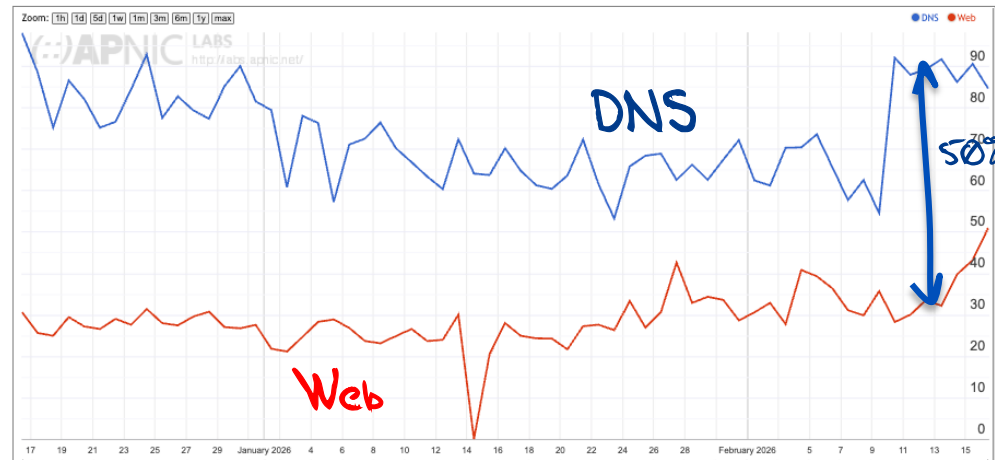
IPv6 DNS Use in Bolivia (BO)



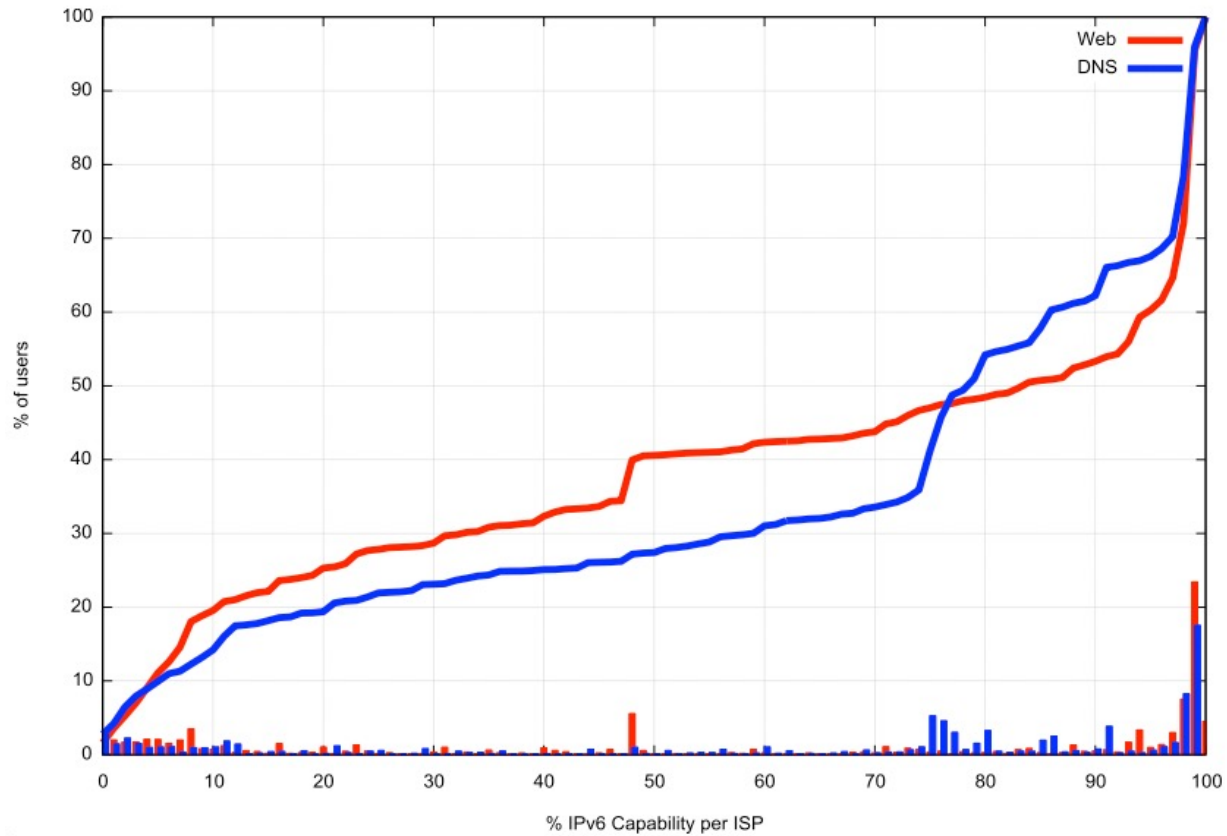
IPv6 DNS Use in Ethiopia (ET)



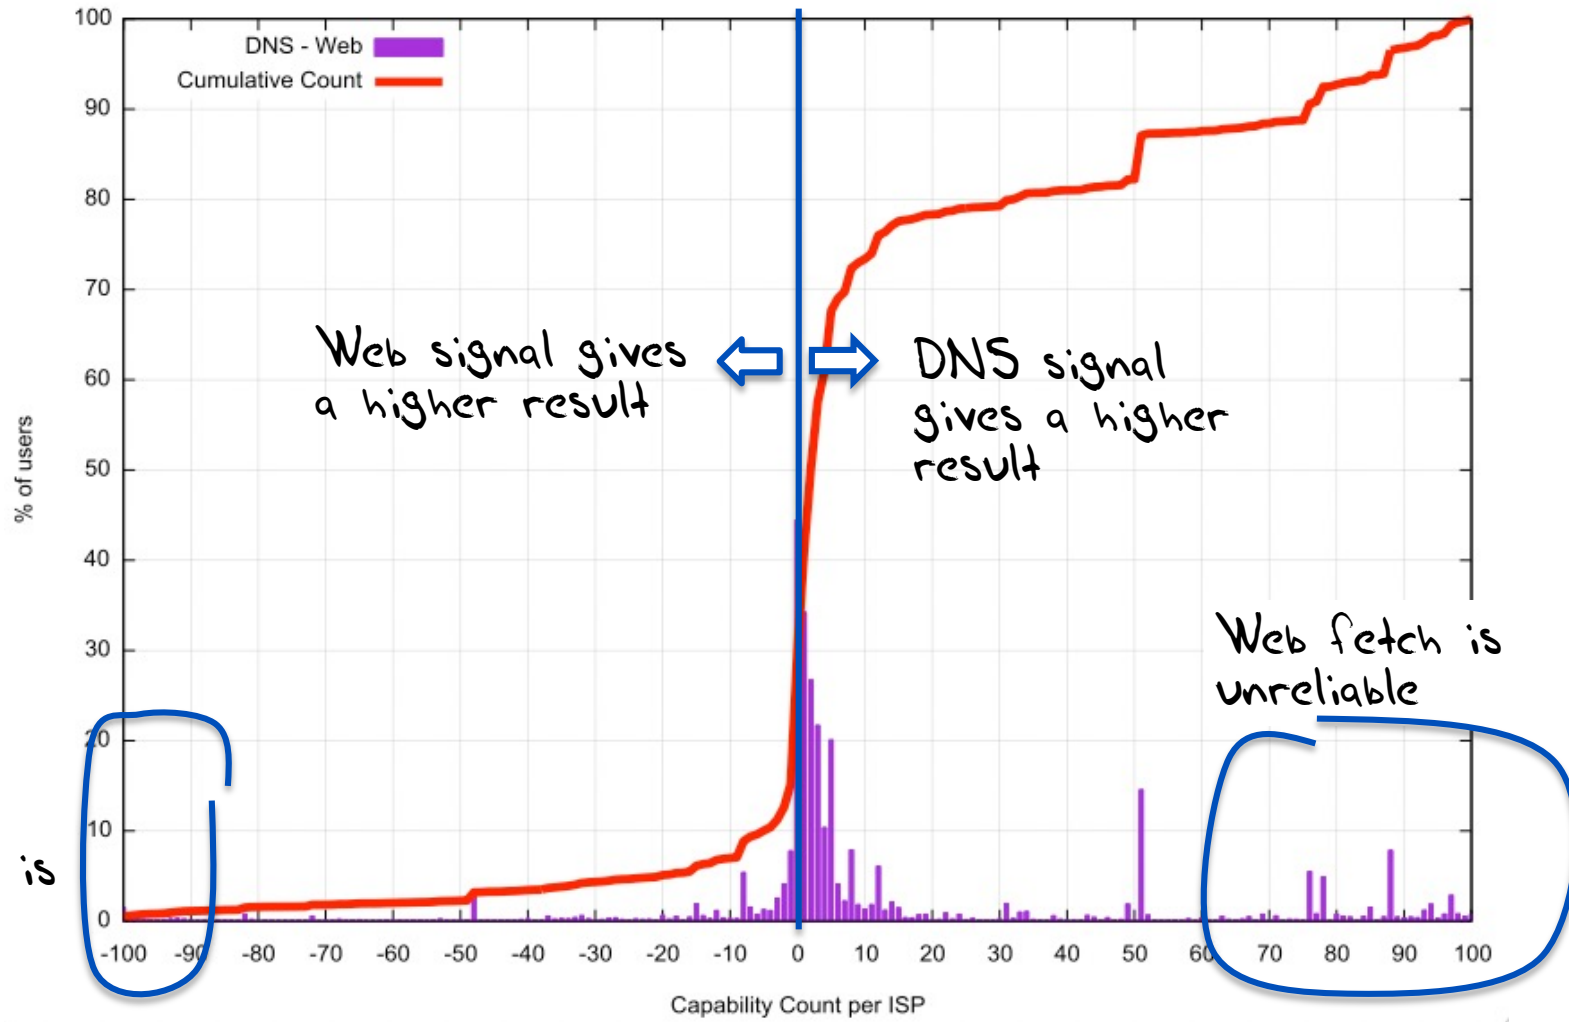
IPv6 DNS Use in Myanmar (MM)



Distribution of Web and DNS Results per Origin AS



DNS vs Web



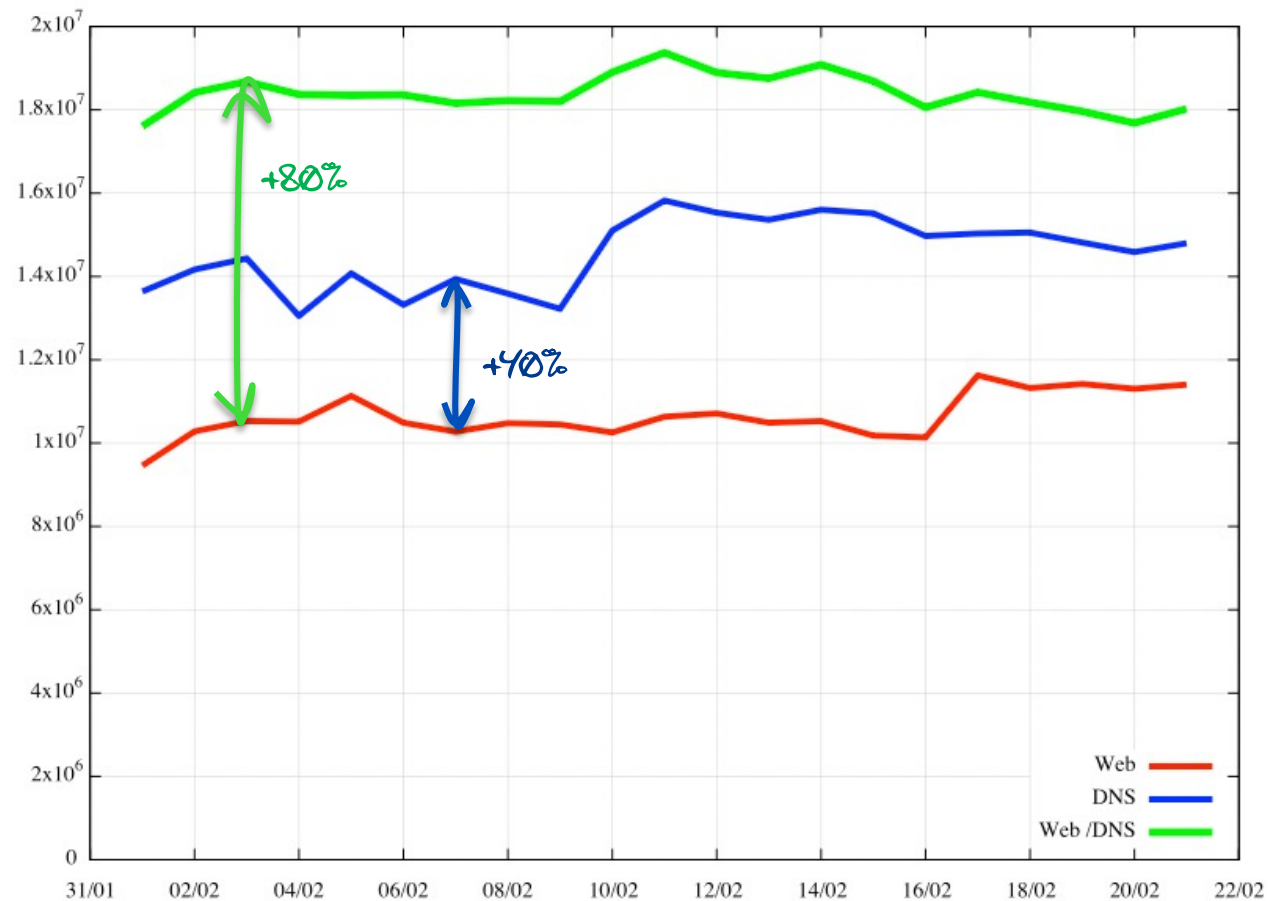
DNS vs Web Measurement

- Here the web results differ from the DNS measurements they appear to be between 10% to 20% lower
- What if we take the maximum of Web and DNS results?

Combined Web and DNS Measurement


Completion Rates:

- The DNS measurement has a significantly higher completion rate compared to the Web-based measurement
- When combined, one third of the completed DNS measurements did not also complete the Web measurement



Combined Web and DNS Measurement

- Adding the Web measurement to the DNS measurement only adds a further 3% to the overall result
- There is a small population of users where the Web fetch indicates a successful query to a IPv6-only server, yet the DNS-only method failed to give an answer
- There is a larger population of users where the DNS-only result indicates success, but there is no web object fetch



Observations

- The deployed DNS behaves in more complex and obscure ways than a simple stub/recursive/authoritative model of DNS infrastructure might suggest
 - Large scale DNS content filtering and DNS monitoring appear to motivate some of this additional DNS behaviour complexity
- Nevertheless, for many DNS behaviour measurements, remote observation of the DNS is more effective using techniques of Glueless Delegation over Web object fetching

That's it!

Questions?