



Measuring IPv6 at web clients and caching resolvers

Emile Aben
emile.aben@ripe.net




Introduction

- We want more insight into IPv6 deployment
- IPv6 deployment numbers:
 - Routing table: 6% of ASes
 - Web traffic: 0.25% - 2 % of web clients
- Where is the difference?
- Measure IPv6 connectivity of end-user combined with ISP infrastructure




Measurement start



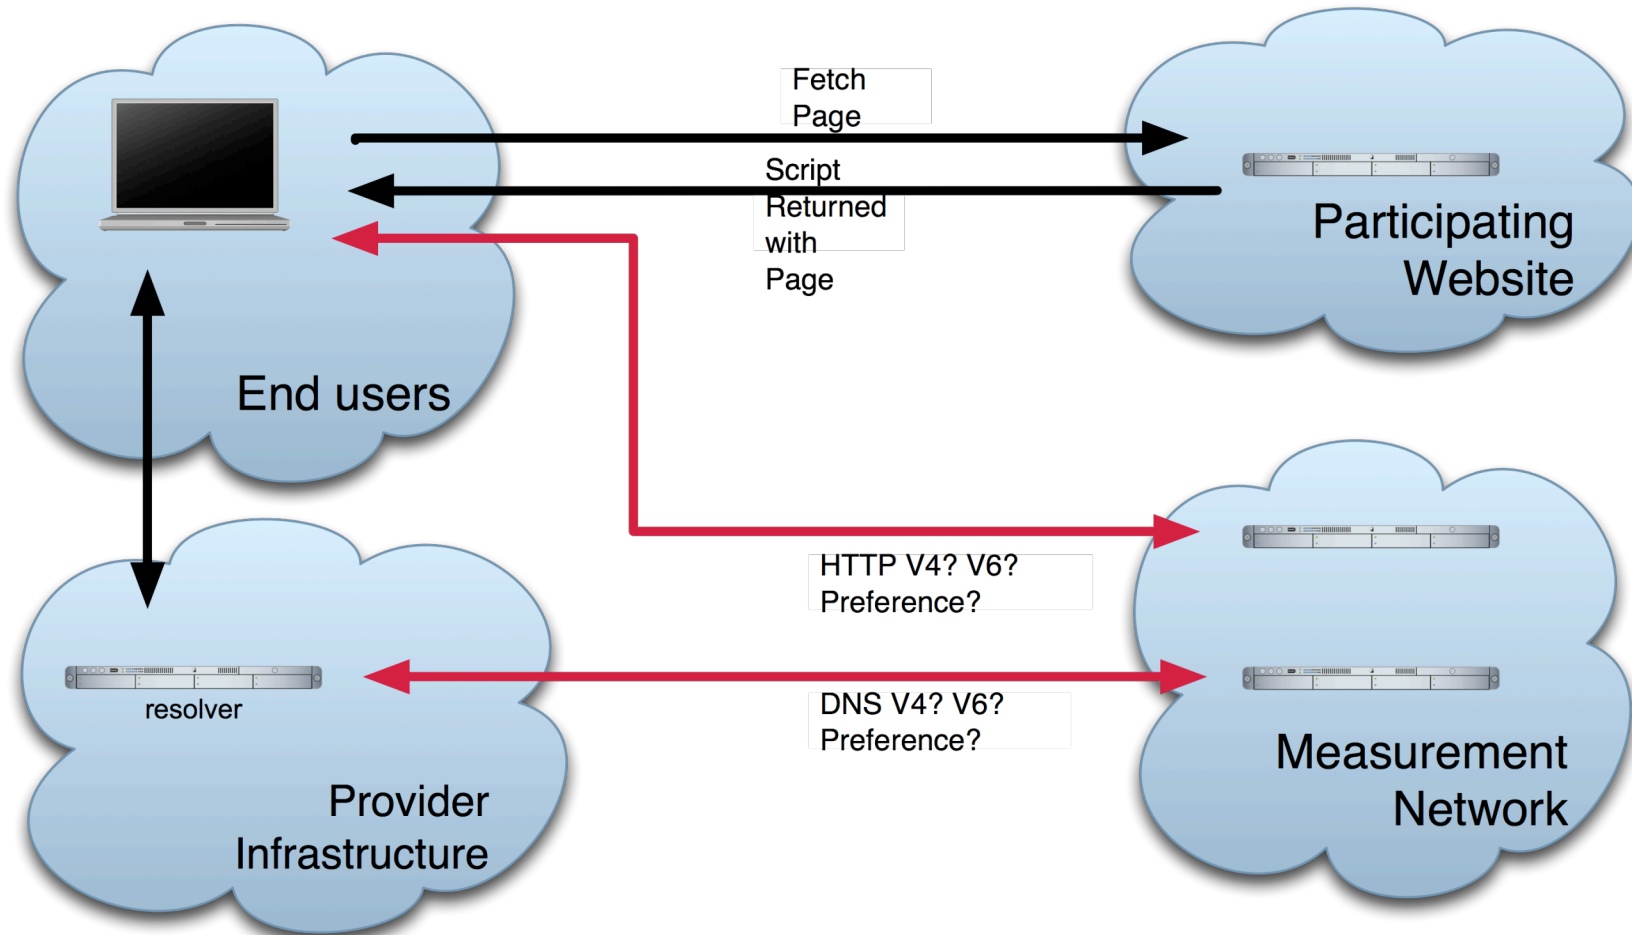


HTTP measurements





DNS measurements





Measurement step 1: javascript


- Web client visits site (www.ripe.net) and fetches piece of javascript
- Script creates a uniqID and causes 4 image lookups:

```
<uniqID>.h4.d6.example.com/img.png?<uniqID>.h4.d6  
<uniqID>.h6.d4.example.com/img.png?<uniqID>.h6.d4  
<uniqID>.hb.db.example.com/img.png?<uniqID>.hb.db  
<uniqID>.h4.d4.example.com/img.png?<uniqID>.h4.d6
```

- h-label: HTTP connectivity (h4,h6,hb)
- d-label: DNS connectivity (d4,d6,db)



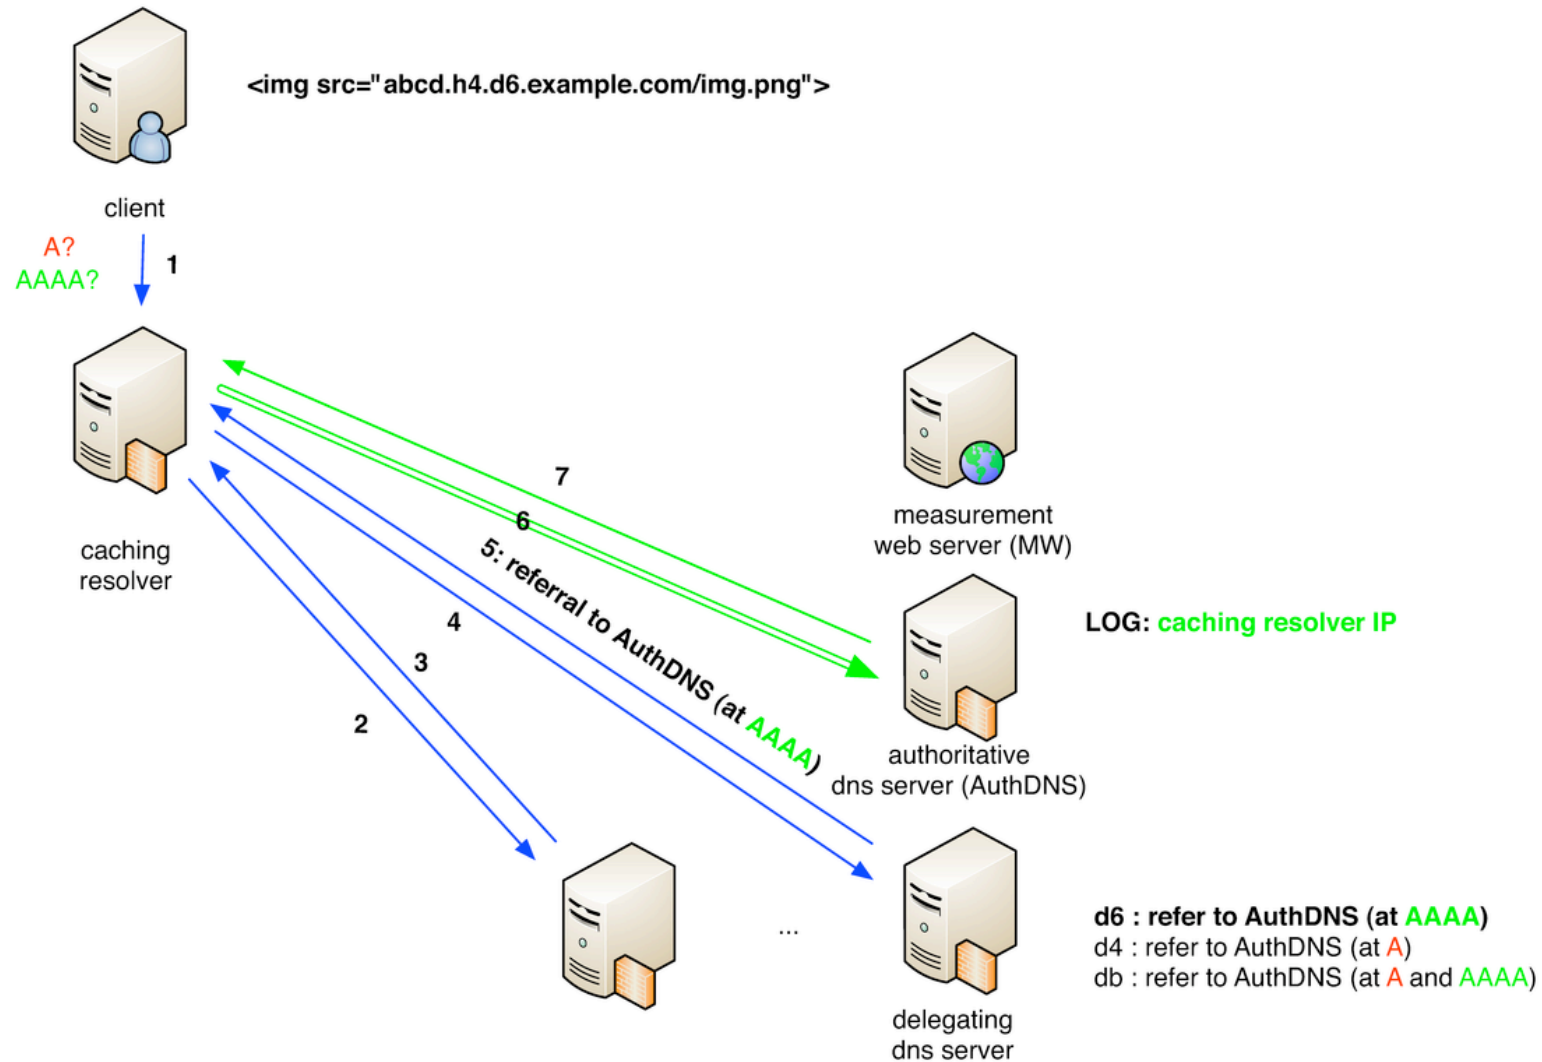
Measurement step 3: HTTP



Authoritative DNS server determines IPv4/IPv6 for client-webserver communication



Measurement step 2: DNS



Delegating DNS server determines IPv4/IPv6 for resolver-authoritative DNS communication



Submeasurements

	DNS v4	DNS v6	DNS both
HTTP v4	h4.d4	h4.d6	-
HTTP v6	h6.d4	-	-
HTTP both	-	-	hb.db



Data collected

HTTP custom logs:

```
1.1.1.1 1273045440 s514725759.h4.d4.example.com "User-agent"
1.1.1.1 1273045440 s514725759.h4.d6.example.com "User-agent"
2010:6:5::1 1273045440 s514725759.hb.db.example.com "User-agent"
2010:6:5::1 1273045440 s514725759.h6.d4.example.com "User-agent"
```

DNS query logs:

```
05-May-2010 09:44:00.531 client 2.2.2.2#30122: view external: query: s514725759.h4.d4.example.com IN A -E
05-May-2010 09:44:00.541 client 2.2.2.2#27174: view external: query: s514725759.h4.d4.example.com IN AAAA -E
05-May-2010 09:44:00.575 client 2.2.2.2#42035: view external: query: s514725759.hb.db.example.com IN A -E
05-May-2010 09:44:00.583 client 2.2.2.2#35884: view external: query: s514725759.hb.db.example.com IN AAAA -E
05-May-2010 09:44:00.593 client 2.2.2.2#1925: view external: query: s514725759.h6.d4.example.com IN A -E
05-May-2010 09:44:00.606 client 2.2.2.2#6064: view external: query: s514725759.h6.d4.example.com IN AAAA -E
05-May-2010 09:44:00.552 client 2010:6:5::2#30403: view external: query: s514725759.h4.d6.example.com IN A -E
05-May-2010 09:44:00.562 client 2010:6:5::2#38536: view external: query: s514725759.h4.d6.example.com IN AAAA -E
```

Mix-and-match:

```
05-May-2010 09:44:00.531 client 2.2.2.2#30122: view external: query: s514725759.h4.d4.example.com IN A -E
05-May-2010 09:44:00.541 client 2.2.2.2#27174: view external: query: s514725759.h4.d4.example.com IN AAAA -E
1.1.1.1 1273045440 s514725759.h4.d4.example.com "User-agent"

05-May-2010 09:44:00.575 client 2.2.2.2#42035: view external: query: s514725759.hb.db.example.com IN A -E
05-May-2010 09:44:00.583 client 2.2.2.2#35884: view external: query: s514725759.hb.db.example.com IN AAAA -E
2010:6:5::1 1273045440 s514725759.hb.db.example.com "User-agent"

05-May-2010 09:44:00.593 client 2.2.2.2#1925: view external: query: s514725759.h6.d4.example.com IN A -E
05-May-2010 09:44:00.606 client 2.2.2.2#6064: view external: query: s514725759.h6.d4.example.com IN AAAA -E
2010:6:5::1 1273045440 s514725759.h6.d4.example.com "User-agent"

05-May-2010 09:44:00.552 client 2010:6:5::2#30403: view external: query: s514725759.h4.d6.example.com IN A -E
05-May-2010 09:44:00.562 client 2010:6:5::2#38536: view external: query: s514725759.h4.d6.example.com IN AAAA -E
1.1.1.1 1273045440 s514725759.h4.d6.example.com "User-agent"
```



Details, details

- UniqID in domain name
 - Allows for correlation of the 4 sub-measurements
 - Forces DNS lookup
- Low DNS RR TTL (just to be sure)
- Measurements are limited to 1 run of the script per day per client
- Local (NCC) traffic filtered out
- Measurement bias on clients:
 - Only visitors to site that hosts measurement-script
 - Clients that use javascript (>95%)




Measurement results



IPv6 for clients/resolvers for www.ripe.net


IPv6 in web clients and the resolvers they use (daily bins)





“Native” IPv6

Native IPv6 in web clients and the resolvers they use (daily bins)



Native: not-autotunneled (ie. No Teredo, 6to4)



Same AS?

- Are things in the same AS?

		n	Same AS	Different AS	Mixed AS
v4 HTTP	v6 HTTP (native)	8k	64%	36%	N/A
v4 HTTP	v4 DNS	520k	79%	19%	2.2%
v6 HTTP	v4 DNS	25k	20%	76%	4.2%
v6 HTTP (native)	v4 DNS	8k	61%	36%	2.4%
v4 HTTP	v6 DNS	31k	61%	38%	0.6%
v4 HTTP	v6 DNS (native)	29k	64%	35%	0.5%

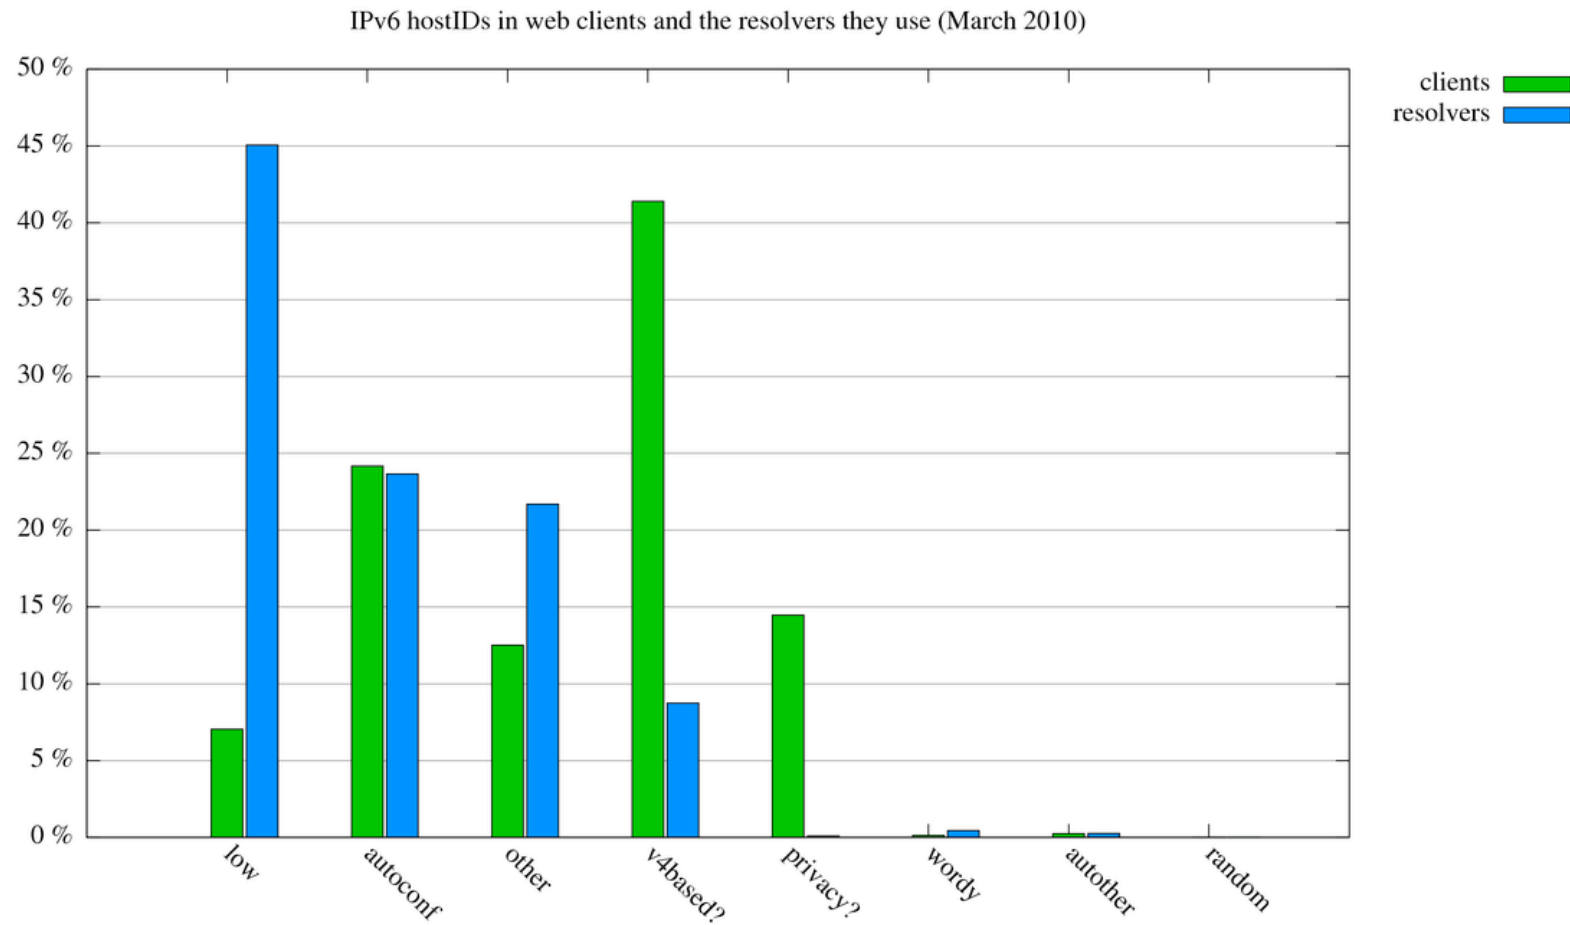


Random facts

- Googlebot does javascript
- In 5% of measurements we see 2 large providers of DNS services cause clientAS != resolverAS
- At least 10% of client v4 AS != client v6 AS caused by ASes involved in tunnel brokering



Host IDs



Classification method: David Malone, PAM 2008



What's next

- Keep this running, we live in interesting times
- We want more data, on Joe Average Internet user
 - You can participate!
 - hosting a piece of javascript on a webpage
 - Ask me =====>
 - Questions, comments?
 - emile.aben@ripe.net





Questions?





Within the AS

CDF of v6/v4 ratio for measurements in ASes where we detect both v4 and native v6 capabilities
(ASes with 50+ hits on v4 or v6)

