Measuring Query Name Minimization

Joao Damas
Geoff Huston

APNIC Labs
October 2020

Quick Summary

NON-query name minimisation resolution sequence

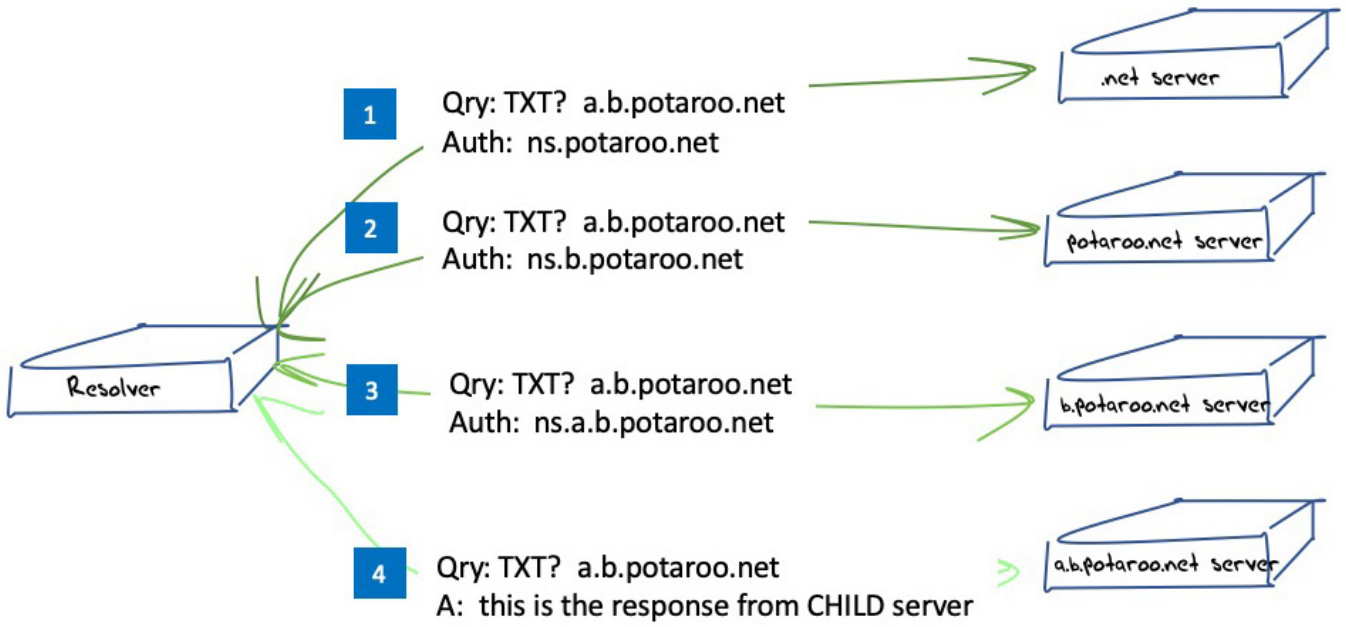

Quick Summary

Query name minimisation technique described in RFC 7816

> Instead of sending the full QNAME and the original QTYPE upstream, a resolver that implements QNAME minimisation and does not already have the answer in its cache sends a request to the name server authoritative for the closest known ancestor of the original QNAME. The request is done with:
>
> o the QTYPE NS
>
> o the QNAME that is the original QNAME, stripped to just one label more than the zone for which the server is authoritative

Quick Summary

Query name minimisation technique described in RFC 7816

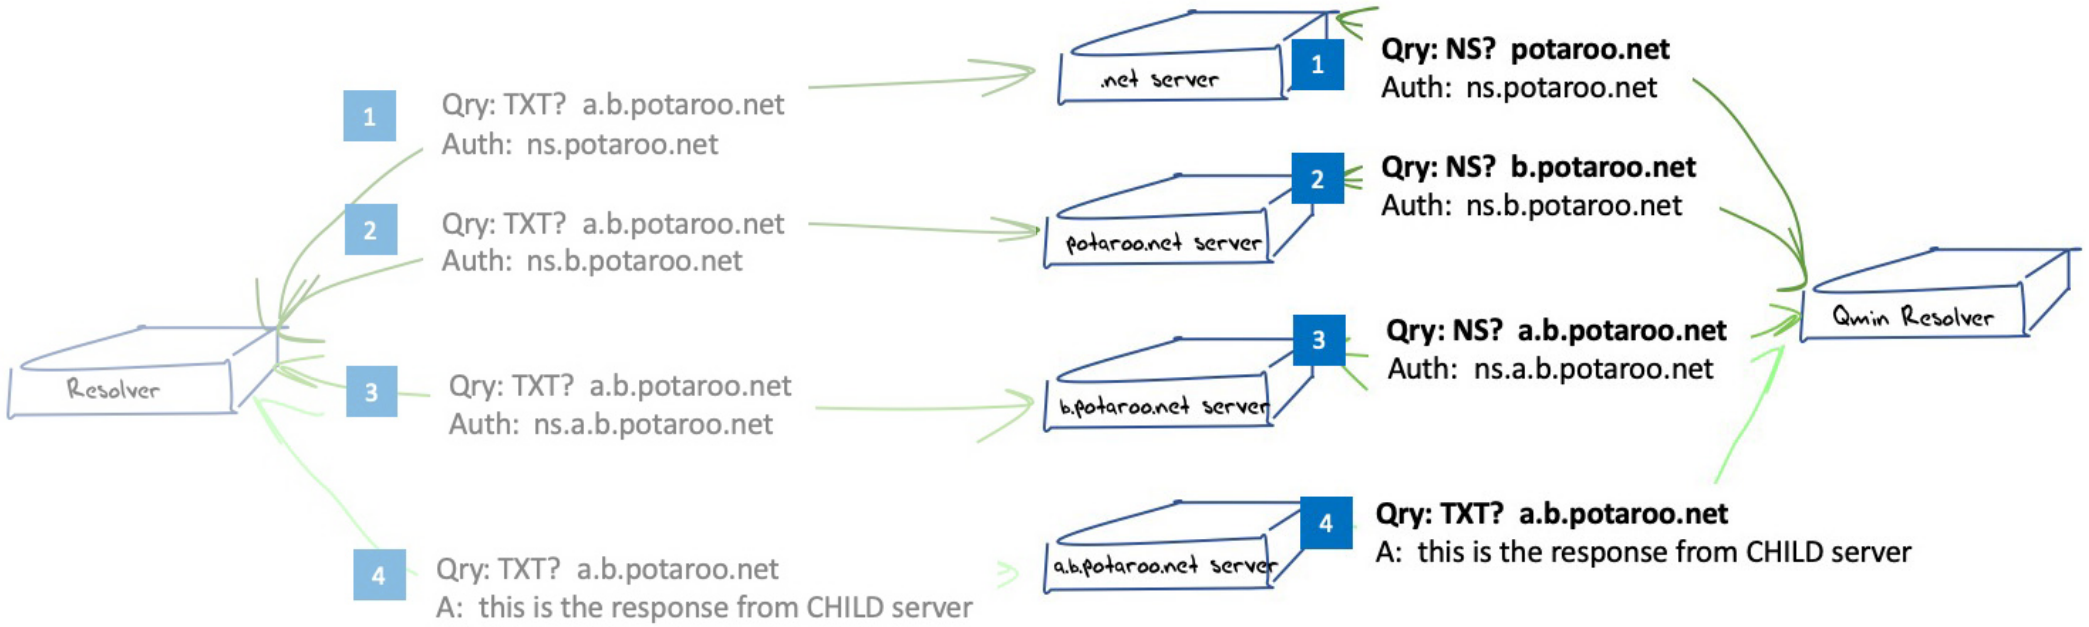

Common Resolver Implementation Status

- BIND 9
  - Implemented in 9.14, active in “relaxed” mode by default
- Unbound
  - Implemented in 1.7.2, active in “non-strict” mode
- Knot
  - Implemented in 1.2.2, active by default
- Power DNS Recursor
  - Implemented in 4.3.0-alpha1, enabled by default since 4.3.0-beta 1

Common Resolver Implementation Status

- BIND 9
  - Implemented in 9.1
- Unb
-
- P

4.3.0-alpha1, enabled by default since 4.3.0-beta 1

it looks like all recursive resolvers that use up-to-date versions of these code bases should be doing query name minimisation by default these days.

What do we see?

Measurement

Let’s look at the adoption of query name minimisation from the perspectives of the end user and their queries, and from the perspective of recursive resolvers

Users whose Queries are handled with Qname Minimization

2019 Results

Experiments	Qmin	Query Type			
		NS	A	AAAA	
429,773,288	11,089,823	2,811,053	8,336,008	1,721	
	3%	1%	2%	0%	% of all experiments
		25%	75%	0%	% of Qmin experiments

Users whose Queries are handled with Qname Minimization

2019 Results

Experiments	Qmin	Query Type			
		NS	A	AAAA	
429,773,288	11,089,823	2,811,053	8,336,008	1,721	
	3%	1%	2%	0%	% of all experiments
		25%	75%	0%	% of Qmin experiments

2020 Results

Experiments	Qmin	Query Type			
		NS	A	AAAA	
357,905,595	63,515,319	4,092,581	59,705,773	-	
	18%	1%	17%	0%	% of all experiments
		6%	94%	0%	% of Qmin experiments

Daily Results - 2020

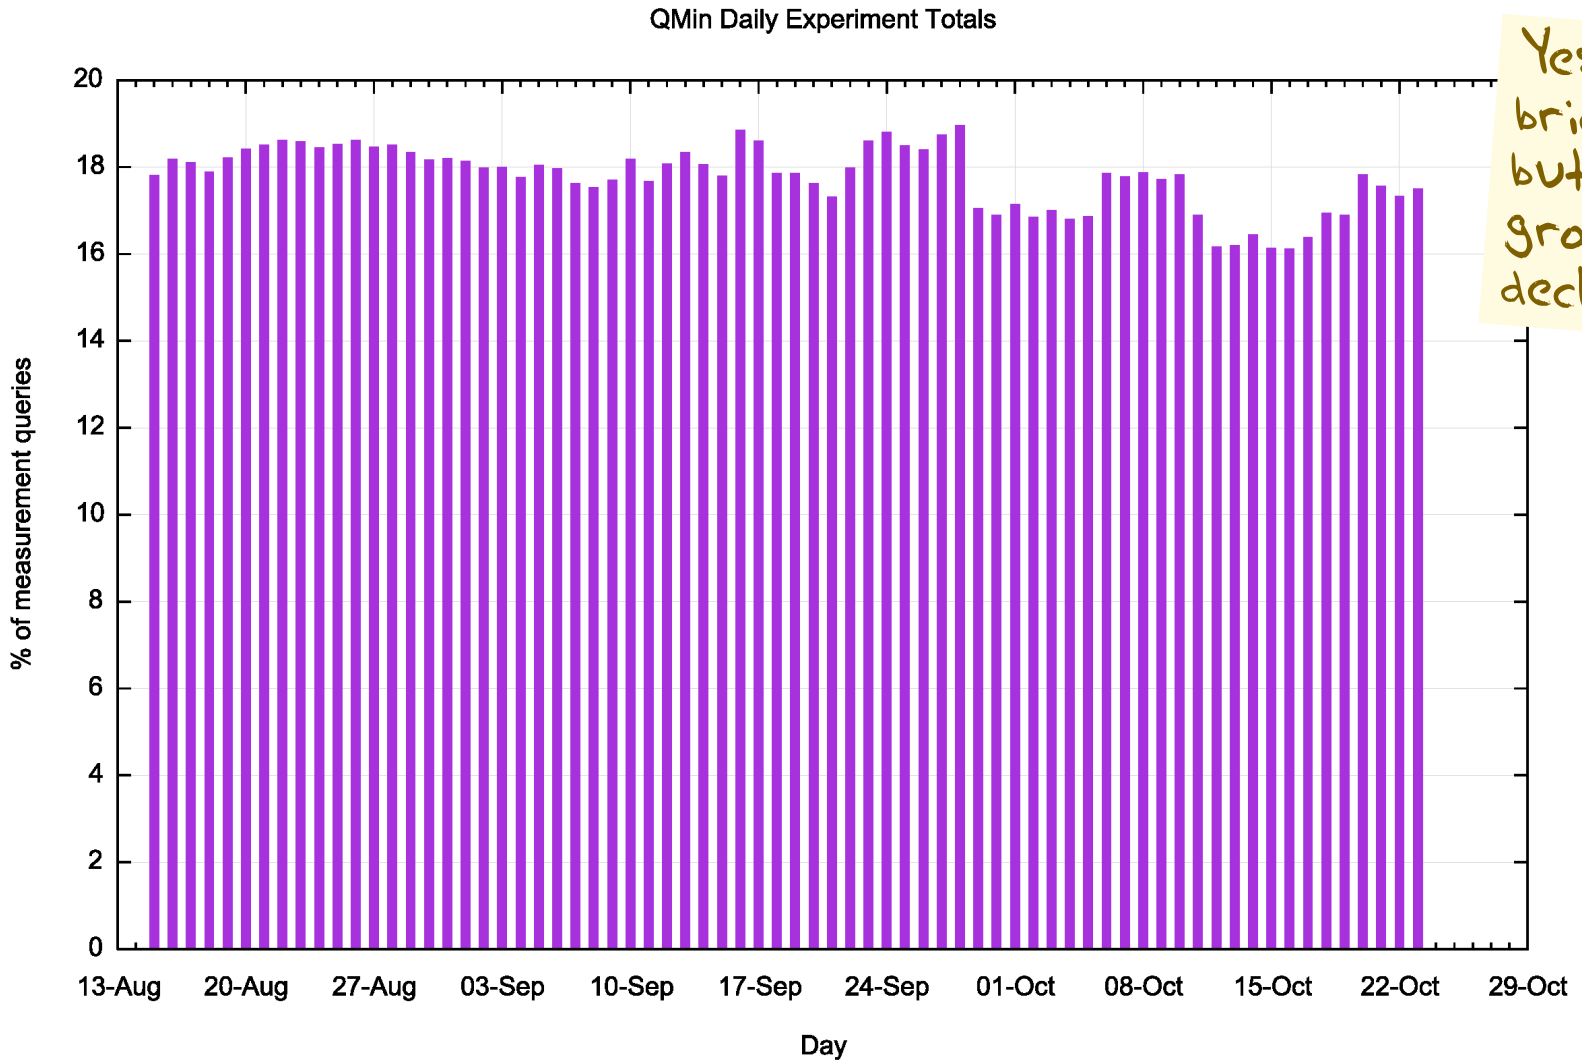

Yes, this is a relatively brief 8-week measurement but the rate is not growing, and may even be declining a little!

Where are these Users?

CC	Qmin	Exps.	Qmin Count	Name
GL	80%	3,433	2,738	Greenland
LI	58%	3,172	1,838	Liechtenstein
MG	56%	423,638	237,652	Madagascar
CY	56%	93,687	52,084	Cyprus
KP	53%	4,192	2,201	DPR Korea
NE	50%	424,271	214,168	Niger
IN	49%	52,608,437	25,665,243	India
GI	48%	3,348	1,616	Gibraltar
NP	48%	634,466	302,691	Nepal
IQ	47%	3,271,159	1,551,627	Iraq
BW	47%	92,113	43,653	Botswana
AF	43%	476,157	205,127	Afghanistan
DE	43%	6,012,110	2,583,028	Germany
IR	41%	5,532,777	2,294,737	Iran
PH	41%	6,384,131	2,605,019	Philippines
SI	39%	151,910	59,964	Slovenia
GE	39%	241,814	94,950	Georgia
TG	39%	121,776	47,704	Togo
MV	38%	33,549	12,658	Maldives
ZW	37%	423,739	158,741	Zimbabwe
GM	36%	47,920	17,031	Gambia
PT	34%	696,889	237,476	Portugal
BY	33%	661,704	220,477	Belarus
ZA	33%	3,084,863	1,022,078	South Africa
NZ	31%	387,654	120,774	New Zealand
FR	30%	4,624,666	1,400,750	France
AD	29%	6,647	1,932	Andorra
GH	29%	1,197,502	346,091	Ghana
MD	29%	293,043	84,263	Moldova
SG	29%	439,993	125,506	Singapore
CM	28%	566,820	161,204	Cameroon
IS	27%	28,563	7,637	Iceland
AO	27%	468,063	124,893	Angola
CG	27%	46,923	12,484	Congo

Resolver Measures

What's a "**resolver**"?

- Always hard to tell these days.
- Over a 16 day period we saw 183,438 distinct IP addresses of resolvers
  - 148,230 IPv4 addresses
    77,548 distinct /24 subnets
  - 35,209 IPv6 addresses
    9,069 distinct /48 subnets

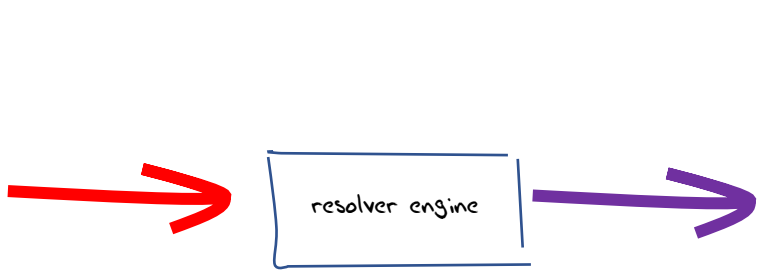

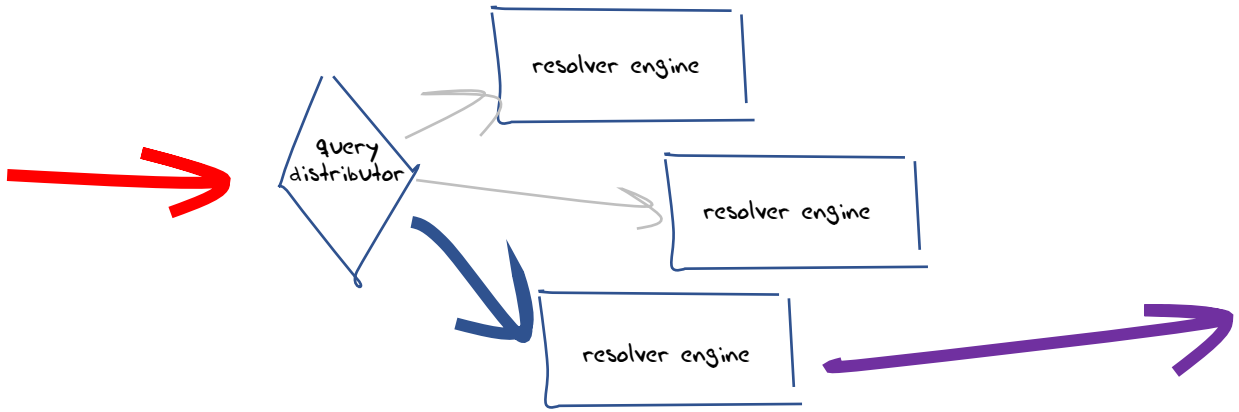

Open Resolvers

Resolver	Qmin Ratio	Experiments	Qmin
googlepdns	0%	222,266,568	2,909
114dns	5%	49,267,636	2,671,180
yandex	0%	28,164,377	238
dnspai	5%	19,787,850	923,698
cloudflare	50%	18,296,672	9,205,045
onedns	7%	15,838,970	1,058,729
opendns	71%	15,488,084	10,997,436
level3	0%	3,083,038	-
quad9	67%	2,537,980	1,703,220
neustar	55%	1,649,393	909,871
vrsgn	0%	1,536,303	-
dyn	55%	558,821	306,645
dnswatch	55%	518,237	287,119
cnnic	0%	515,878	-
greenteamdns	0%	421,532	114
he	83%	176,262	146,637
comodo	26%	112,308	29,613
freedns	0%	87,804	-
dnspod	0%	54,164	46

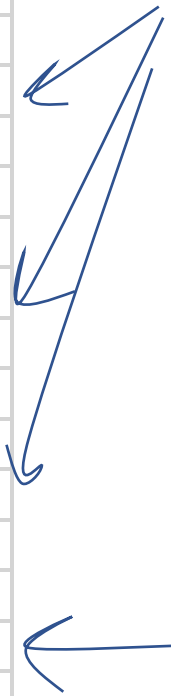

What's behind these 50+-70+ ratios? is Qmin only partially deployed in the DNS service anycast constellation?

This is more expected!

ISP Resolvers

ASN	QMin Ratio	Experiments	Qmin	Name	CC
4134	8%	272,985,533	22,389,630	CHINANET-BACKBONE	CN
55836	56%	103,846,458	58,615,952	Reliance Jio	IN
4837	5%	52,525,073	2,884,098	CHINA UNICOM	CN
9808	5%	44,902,506	2,399,098	Guangdong Mobile	CN
9498	0%	36,424,784	113	BHARTI Airtel BBIL	IN
58543	0%	35,255,383	-	CHINATELECOM Guangdong	CN
56046	41%	31,490,572	12,941,229	China Mobile Jiangsu	CN
56040	0%	19,782,214	144	China Mobile Guangdong	CN
7922	0%	18,081,958	2,460	COMCAST	US
4835	47%	15,634,509	7,345,689	CHINANET-IDC-SN China Telecom	CN
24560	0%	14,859,198	62	Bharti Airtel Broadband	IN
56041	0%	10,645,009	48,689	China Mobile Zhejiang	CN
6730	50%	9,398,245	4,723,646	SUNRISE	CH
24445	1%	8,922,489	85,080	Henan Mobile	CN
38266	1%	8,895,802	125,353	Vodafone India	IN
7552	0%	8,891,315	636	Viettel	VN
17676	2%	8,714,412	199,840	Softbank BB	JP
30986	32%	8,029,250	2,546,706	SCANCOM	GH
8151	0%	7,881,161	1,193	Uninet	MX
7018	0%	7,870,637	874	ATT INTERNET	US
28573	0%	7,837,132	521	CLARO	BR
4766	0%	7,629,352	280	Korea Telecom	KR
9121	0%	7,340,736	826	TTNET	TR
27725	0%	6,661,765	12,907	Empresa de Telecomunicaciones de Cuba	CU
3462	0%	6,599,708	452	HINET	TW

Observations

- Query name minimisation is gathering momentum in the past 12 months (3% or users in mid 2019 to 18% of users in mid-2020)
- While all common vendor code has enabled Query name minimisation, enabling this behaviour in ISP and open resolvers is fragmentary
  - Why is it not deployed? What's the concern?

Our Measurement

- We are using the 4th and 5th level names to perform the experiment

  <*unique-label*> **.** ent-<*unique label*> **.** <geo-code> **.** <common_name> **.** net

  - Some resolvers (Google?) only perform Qname minimisation to the 3rd level
  - Why?

  - **Is privacy no longer important at the bottom of the name hierarchy?**
  - **Or is it only TLD servers that breach privacy in query names?**
  - **Or are recursive operators just making it up on the fly?**

More Questions

- **Where and why is Query Name minimisation important?**
- Does it differ by scale?
  - Small scale recursive resolvers at the edge of the network?
  - ISP-operated recursive resolvers?
  - Open recursive resolvers?

- Is the query name alone a privacy threat or is the combination of the recursive resolver with the query name the problem?

Last Question

What's the most critical privacy risk in today's DNS?

- ❑ Explicit Client Subnet?
- ❑ Full query name without attribution from recursive to authoritative?
- ❑ Recursive resolvers seeing both the full query name and attribution?
- ❑ Unencrypted stub-to-recursive DNS transactions?
- ❑ Unencrypted recursive-to-authoritative DNS transactions?

Thanks!