

Measuring Query Name Minimization

Geoff Huston

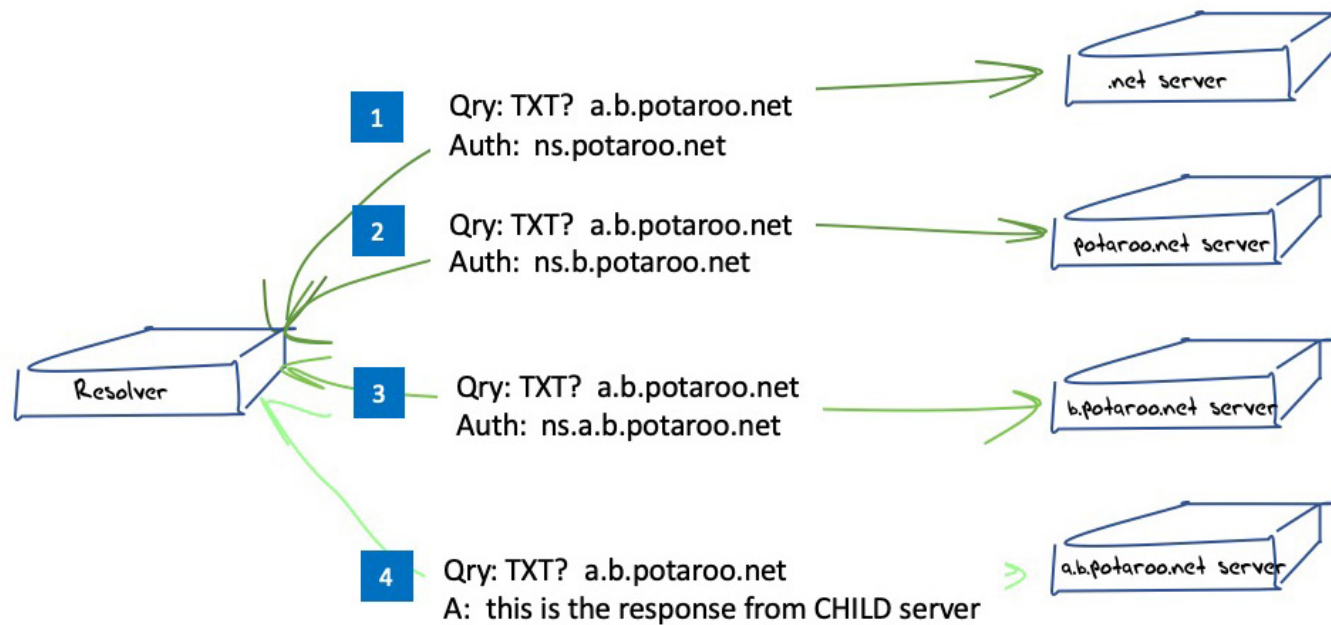
Joao Damas

APNIC Labs

September 2020

Quick Summary

NON-query name minimisation resolution sequence



Quick Summary

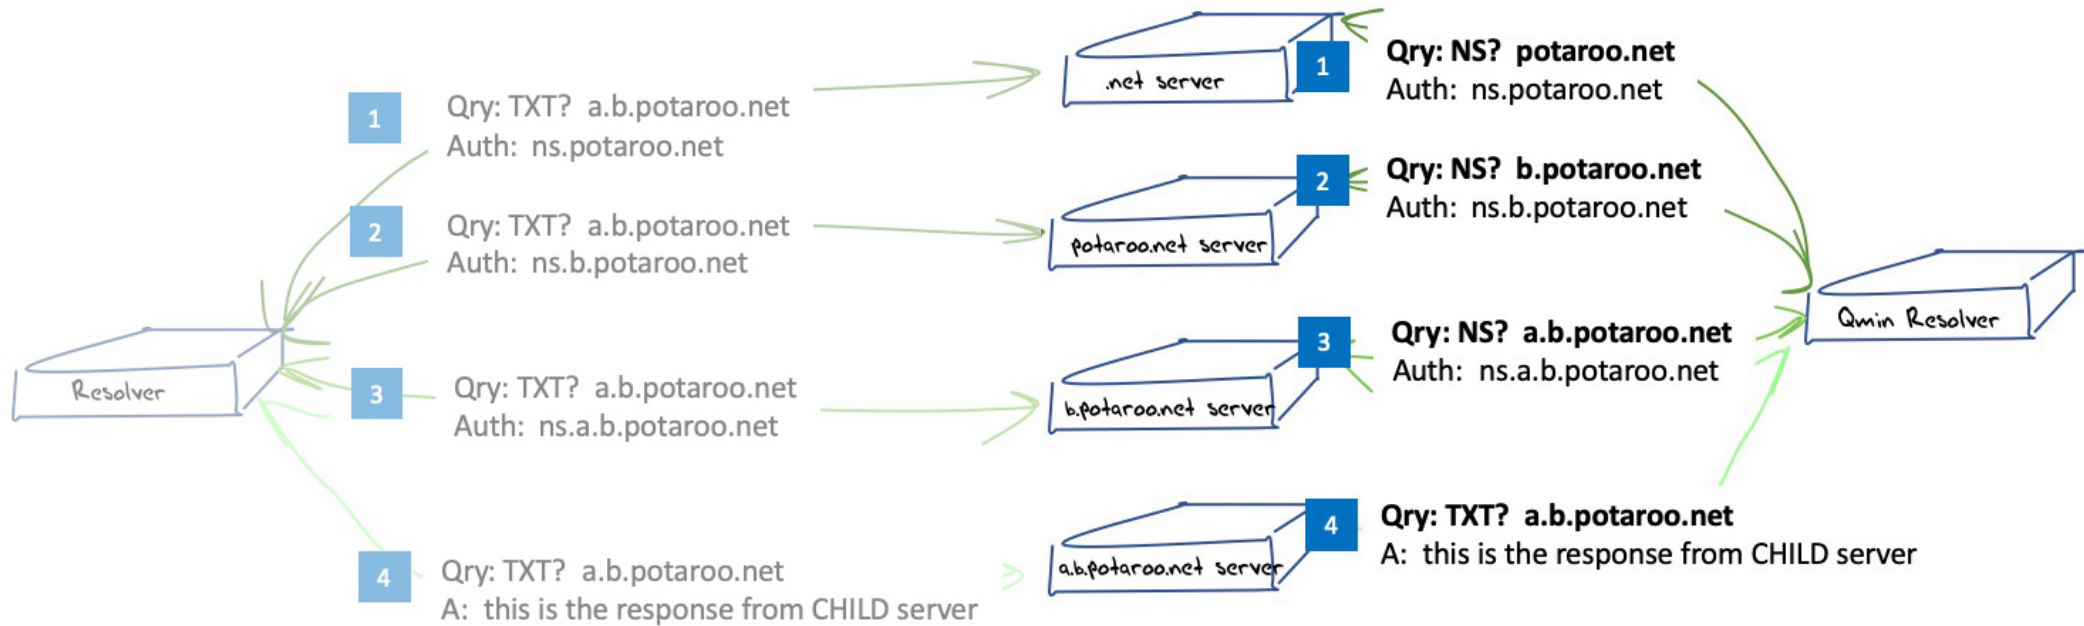
Query name minimisation technique described in RFC 7816

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 9.14, active in “rel-”
- Unbound
 - Implemented in 4.3.0-alpha1, enabled by default since 4.3.0-beta 1
- P
Right? `resolv`
 - Implemented in 4.3.0-alpha1, enabled by default since 4.3.0-beta 1
- It looks like all recursive resolvers should be doing query name minimisation these days.

Measurements

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

User Measurements

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

User Measurements

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	
99,303,954	18,219,251	1,411,485	16,880,583	0	
	18%	1%	16%	0%	% of all experiments
		7%	92%	0%	% of Qmin experiments

User Measures

The proportion of users who use recursive resolvers that perform Query Name minimization has risen from 3% of users to 18% of users in the past 12 months.

The common resolver behaviour is to perform the discovery queries using query type A, not NS or AAAA

Where are these Users?

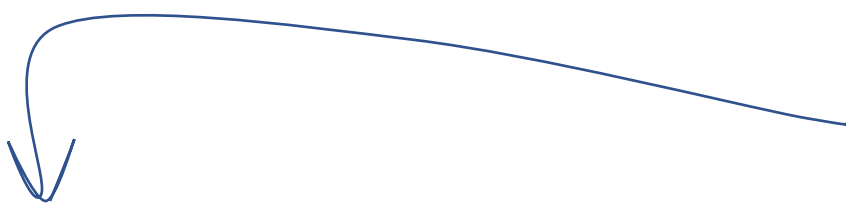
CC	Qmin Ratio	Experiments	Qmin	CC Name
AD	60%	1856	1128	Andorra
CY	59%	26,011	15,380	Cyprus
IR	57%	1,525,556	876,474	Iran
BW	56%	25,598	14,503	Botswana
NE	56%	114782	64708	Niger
KP	53%	1149	613	DPR Korea
IN	52%	14496031	7606073	India
NP	51%	175589	91016	Nepal
MV	44%	9303	4179	Maldives
ZW	42%	117058	49396	Zimbabwe
AF	41%	131402	54505	Afghanistan
GM	41%	13273	5460	Gambia
PT	39%	193912	76982	Portugal
DE	37%	1681871	626608	Germany
MG	36%	114996	41848	Madagascar
BM	36%	1462	532	Bermuda
BY	35%	183944	64963	Belarus
FR	35%	1290872	453342	France
ZA	34%	858408	297725	South Africa
NG	34%	1540623	532479	Nigeria
IQ	32%	886273	291240	Iraq
GE	32%	66979	21725	Georgia
NZ	31%	108164	34398	New Zealand
SI	31%	42269	13434	Slovenia
CG	30%	13015	3911	Congo
EC	29%	275295	80786	Ecuador

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



Open DNS re	Qmin Ratio	Experiments	Qmin Experiments
googlepdns	0%	70,253,285	530
114dns	6%	6,590,007	370,145
cloudflare	50%	6,178,049	3,104,221
opendns	68%	5,490,013	3,717,390
dnspai	4%	4,700,366	165,689
onedns	10%	3,175,244	317,439
vrsgn	0%	927,125	0
quad9	70%	760,514	534,262
level3	0%	407,443	0
neustar	59%	394,491	231,720
yandex	0%	246,648	0
dnswatch	56%	167,079	94,226
dyn	58%	109,400	63,938
cnnic	0%	54,841	0
he	98%	39,240	38,311

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 Experiments	AS Name	
55836	59%	28,468,311	16,687,545	RELIANCEJIO-IN Reliance Jio Infocomm Limited	IN
4837	5%	11,328,596	622,812	CHINA169-BACKBONE CHINA UNICOM China169 Backbone	CN
9808	8%	9,999,271	824,254	CMNET-GD Guangdong Mobile Communication Co.Ltd.	CN
9498	0%	8,632,305	5	BBIL-AP BHARTI Airtel Ltd.	IN
58543	0%	8,579,142	0	CHINATELECOM-GUANGDONG-IDC Guangdong	CN
56046	51%	6,935,852	3,561,604	CMNET-JIANGSU-AP China Mobile communications corporation	CN
7922	0%	5,590,304	620	COMCAST-7922	US
6730	50%	4,714,894	2,359,887	SUNRISE	CH
24560	0%	3,467,579	4	AIRTELBBROADBAND-AS-AP Bharti Airtel Ltd.	IN
4835	51%	3,380,595	1,729,101	CHINANET-IDC-SN China Telecom (Group)	CN
30986	32%	3,170,873	1,016,619	SCANCOM	GH
56040	0%	3,163,238	0	CMNET-GUANGDONG-AP China Mobile communications corpor	CN
28573	0%	2,837,684	123	CLARO S.A.	BR
7018	0%	2,457,623	185	ATT-INTERNET4	US
7552	0%	2,242,776	81	VIETEL-AS-AP Viettel Group	VN
8151	0%	2,121,346	277	Uninet S.A. de C.V.	MX
22394	0%	1,956,376	0	CELLCO	US
12322	60%	1,795,299	1,074,475	PROXAD	FR
38266	1%	1,737,413	22,059	VODAFONE-IN Vodafone India Ltd.	IN
17676	2%	1,721,203	37,873	GIGAINFRA Softbank BB Corp.	JP
56041	1%	1,654,582	12,526	CMNET-ZHEJIANG-AP China Mobile communications corporatio	CN
17633	64%	1,611,910	1,030,195	CHINATELECOM-SD-AS-AP ASN for Shandong Provincial Net of C	CN
24445	1%	1,520,748	11,555	CMNET-V4HENAN-AS-AP Henan Mobile Communications Co.	CN
17799	0%	1,431,372	0	CHINATELECOM-LN-AS-AP asn for Liaoning Provincial Net of CT	CN
24444	0%	1,382,238	0	CMNET-V4SHANDONG-AS-AP Shandong Mobile Communicator	CN

Observations

- Query name minimisation is gathering momentum in the past 12 months (3% of 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?

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?
- Are there residual issues with handling of empty non-terminals?

Thanks!