

DNS query name minimisation to improve privacy
draft-ietf-dnsop-qname-minimisation-01

Abstract

This document describes one of the techniques that could be used to improve DNS privacy (see [[I-D.ietf-dprive-problem-statement](#)]), a technique called "qname minimisation".

Discussions of the document should take place on the DNSOP working group mailing list [[dnsop](#)].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 19, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in [Section 4](#).e of

the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction and background	2
2.	Qname minimisation	2
3.	Operational considerations	3
4.	Performance implications	5
5.	Security considerations	5
6.	Acknowledgments	6
7.	References	6
7.1.	Normative References	6
7.2.	Informative References	6
7.3.	URIs	7
Appendix A.	An algorithm to find the zone cut	7
	Author's Address	8

[1.](#) Introduction and background

The problem statement is exposed in [\[I-D.ietf-dprive-problem-statement\]](#) TODO: add a reference to the specific section when ietf-dprive-problem-statement will be published as RFC. The terminology ("qname", "resolver", etc) is also defined in this companion document. This specific solution is not intended to fully solve the DNS privacy problem; instead, it should be viewed as one tool amongst many.

It follows the principle explained in [section 6.1 of \[RFC6973\]](#): the less data you send out, the less privacy problems you'll get.

[2.](#) Qname minimisation

The idea is to minimise the amount of data sent from the DNS resolver. When a resolver receives the query "What is the AAAA record for www.example.com?", it sends to the root (assuming a cold resolver, whose cache is empty) the very same question. Sending "What are the NS records for .com?" would be sufficient (since it will be the answer from the root anyway). To do so would be compatible with the current DNS system and therefore could be easily deployable, since it is an unilateral change to the resolvers, it does not change the protocol. Because of that, resolver implementers may do qname minimisation in slightly different ways.

If "minimisation" is too long, you can write it "m10n".

To do such minimisation, the resolver needs to know the zone cut [\[RFC2181\]](#). Zone cuts do not necessarily exist at every label

boundary. If we take the name `www.foo.bar.example`, it is possible that there is a zone cut between "foo" and "bar" but not between "bar" and "example". So, assuming the resolver already knows the name servers of `.example`, when it receives the query "What is the AAAA record of `www.foo.bar.example`", it does not always know if the request should be sent to the name servers of `bar.example` or to those of `example`. [RFC2181] suggests a method to find the zone cut (section 6), so resolvers may try it.

Note that DNSSEC-validating resolvers already have access to this information, since they have to find the zone cut (the DNSKEY record set is just below, the DS record set just above).

Minimising the amount of data sent also, in part, addresses the case of a wire sniffer as well the case of privacy invasion by the servers.

One should note that the behaviour suggested here (minimising the amount of data sent in qnames) is NOT forbidden by the [RFC1034] (section 5.3.3) or [RFC1035] (section 7.2). Sending the full qname to the authoritative name server is a tradition, not a protocol requirement. This tradition comes[*mockapetris-history*] from a desire to optimize the number of requests, when the same name server is authoritative for many zones in a given name (something which was more common in the old days, where the same name servers served `.com` and the root) or when the same name server is both recursive and authoritative (something which is strongly discouraged now). Whatever the merits of this choice at this time, the DNS is quite different now.

As mentioned before, there are several ways to implement qname minimisation. Two main strategies are the aggressive one and the lazy one. In the aggressive one, the resolver only sends NS queries as long as it does not know the zone cuts. This is the safest, from a privacy point of view. The lazy way "piggybacks" on the traditional resolution code. It sends traditional full qnames and learn the zone cuts from the referrals received, then switching to NS queries. This leaks more data but probably requires less changes in the existing resolver codebase.

3. Operational considerations

The administrators of the forwarders, and of the authoritative name servers, will get less data, which will reduce the utility of the statistics they can produce (such as the percentage of the various qtypes). On the other hand, it may decrease their legal responsibility.

Some broken name servers do not react properly to qtype=NS requests. For instance, some authoritative name servers embedded in load balancers reply properly to A queries but send REFUSED to NS queries. REMOVE THIS SENTENCE BEFORE PUBLICATION: As an example of today, look at www.ratp.fr (not ratp.fr). This behaviour is a gross protocol violation, and there is no need to stop improving the DNS because of such brokenness. However, qname minimisation may still work with such domains since they are only leaf domains (no need to send them NS requests). Such setup breaks more than just qname minimisation. It breaks negative answers, since the servers don't return the correct SOA, and it also breaks anything dependent upon NS and SOA records existing at the top of the zone.

A problem can also appear when a name server does not react properly to ENT (Empty Non-Terminals). If ent.example.com has no resource records but foobar.ent.example.com does, then ent.example.com is an ENT. A query, whatever the qtype, for ent.example.com must return NODATA (NOERROR / ANSWER: 0). However, some broken name servers return NXDOMAIN for ENTs. REMOVE THIS SENTENCE BEFORE PUBLICATION: As an example of today, look at com.akadns.net or www.upenn.edu with its delegations to Akamai. If a resolver queries only foobar.ent.example.com, everything will be OK but, if it implements qname minimisation, it may query ent.example.com and get a NXDOMAIN. See also section 3 of [[I-D.vixie-dnsexst-resimprove](#)] for the other bad consequences of this brokenness.

Another way to deal with such broken name servers would be to try with A requests (A being chosen because it is the most common and hence a qtype which will be always accepted, while a qtype NS may ruffle the feathers of some middleboxes). Instead of querying name servers with a query "NS example.com", we could use "A [_](#).example.com" and see if we get a referral.

Other strange and illegal practices may pose a problem: for instance, there is a common DNS anti-pattern used by low-end web hosters that also do DNS hosting that exploits the fact that the DNS protocol (pre-DNSSEC) allows certain serious misconfigurations, such as parent and child zones disagreeing on the location of a zone cut. Basically, they have a single zone with wildcards like:

```
*.example.          60  IN  A    192.0.2.6
```

(It is not known why they don't just wildcard all of "*" and be done with it.)

This lets them turn up tons of web hosting customers without having to configure thousands of individual zones on their nameservers. They just tell the prospective customer to point their NS records at

their nameservers, and the Web hoster doesn't have to provision anything in order to make the customer's domain resolve.

Qname minimisation can decrease performance in some cases, for instance for a deep domain name (like `www.host.group.department.example.com` where `host.group.department.example.com` is hosted on `example.com`'s name servers). For such a name, a cold resolver will, depending how qname minimisation is implemented, send more queries. Once warm, there will be no difference with a traditional resolver. A possible solution is to always use the traditional algorithm when the cache is cold and then to move to qname minimisation. This will decrease the privacy a bit but will guarantee no degradation of performance.

Another useful optimisation may be, in the spirit of the HAMMER idea [[I-D.wkumari-dnsop-hammer](#)] to probe in advance for the introduction of zone cuts where none previously existed (i.e. confirm their continued absence, or discover them.)

4. Performance implications

The main goal of qname minimisation is to improve privacy by sending less data. However, it may have other advantages. For instance, if a root name server receives a query from some resolver for `A.CORP` followed by `B.CORP` followed by `C.CORP`, the result will be three `NXDOMAINS`, since `.CORP` does not exist in the root zone. Under query name minimisation, the root name servers would hear only one question (for `.CORP` itself) to which they could answer `NXDOMAIN`, thus opening up a negative caching opportunity in which the full resolver could know a priori that neither `B.CORP` or `C.CORP` could exist. Thus in this common case the total number of upstream queries under qname minimisation would be counter-intuitively less than the number of queries under the traditional iteration (as described in the DNS standard).

Qname minimisation may also improve look-up performance for TLD operators. For a typical TLD, delegation-only, and with delegations just under the TLD, a 2-label QNAME query is optimal for finding the delegation owner name.

5. Security considerations

No security consequence (besides privacy improvment) is known at this time.

6. Acknowledgments

Thanks to Olaf Kolkman for the original idea although the concept is probably much older [1]. Thanks to Mark Andrews and Francis Dupont for the interesting discussions. Thanks to Brian Dickson, Warren Kumari and David Conrad for remarks and suggestions. Thanks to Mohsen Souissi for proofreading. Thanks to Tony Finch for the zone cut algorithm in [Appendix A](#). Thanks to Paul Vixie for pointing out that there are practical advantages (besides privacy) to qname m10n. Thanks to Phillip Hallam-Baker for the fallback on A queries, to deal with broken servers. Thanks to Robert Edmonds for an interesting anti-pattern.

7. References

7.1. Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC6973] Cooper, A., Tschofenig, H., Aboba, B., Peterson, J., Morris, J., Hansen, M., and R. Smith, "Privacy Considerations for Internet Protocols", [RFC 6973](#), July 2013.
- [I-D.ietf-dprive-problem-statement] Bortzmeyer, S., "DNS privacy considerations", [draft-ietf-dprive-problem-statement-01](#) (work in progress), January 2015.

7.2. Informative References

- [RFC2181] Elz, R. and R. Bush, "Clarifications to the DNS Specification", [RFC 2181](#), July 1997.
- [I-D.wkumari-dnsop-hammer] Kumari, W., Arends, R., Woolf, S., and D. Migault, "Highly Automated Method for Maintaining Expiring Records", [draft-wkumari-dnsop-hammer-01](#) (work in progress), July 2014.
- [I-D.vixie-dnsext-resimprove] Vixie, P., Joffe, R., and F. Neves, "Improvements to DNS Resolvers for Resiliency, Robustness, and Responsiveness", [draft-vixie-dnsext-resimprove-00](#) (work in progress), June 2010.

[dnsop] IETF, , "The DNSOP working group of IETF", March 2014,
<<https://datatracker.ietf.org/wg/dnsop/charter/>>.

[mockapetris-history]
Mockapetris, P., "Private discussion", January 2015.

7.3. URIs

[1] <https://lists.dns-oarc.net/pipermail/dns-operations/2010-February/005003.html>

Appendix A. An algorithm to find the zone cut

Although a validating resolver already has the logic to find the zone cut, other resolvers may be interested by this algorithm to follow in order to locate this cut:

(0) If the query can be answered from the cache, do so, otherwise iterate as follows:

(1) Find closest enclosing NS RRset in your cache. The owner of this NS RRset will be a suffix of the QNAME - the longest suffix of any NS RRset in the cache. Call this PARENT.

(2) Initialize CHILD to the same as PARENT.

(3) If CHILD is the same as the QNAME, resolve the original query using PARENT's name servers, and finish.

(4) Otherwise, add a label from the QNAME to the start of CHILD.

(5) If you have a negative cache entry for the NS RRset at CHILD, go back to step 3.

(6) Query for CHILD IN NS using PARENT's name servers. The response can be:

(6a) A referral. Cache the NS RRset from the authority section and go back to step 1.

(6b) An authoritative answer. Cache the NS RRset from the answer section and go back to step 1.

(6c) An NXDOMAIN answer. Return an NXDOMAIN answer in response to the original query and stop.

(6d) A NOERROR/NODATA answer. Cache this negative answer and go back to step 3.

Author's Address

Stephane Bortzmeyer
AFNIC
1, rue Stephenson
Montigny-le-Bretonneux 78180
France

Phone: +33 1 39 30 83 46
Email: bortzmeyer+ietf@nic.fr
URI: <http://www.afnic.fr/>