

Domain Name System Operations (dnsop) Working Group
Internet-Draft
Updates: [1034](#),2308 (if approved)
Intended status: Standards Track
Expires: June 1, 2016

S. Bortzmeyer
AFNIC
S. Huque
Verisign labs
November 29, 2015

NXDOMAIN really means there is nothing underneath
draft-bortzmeyer-dnsop-nxdomain-cut-02

Abstract

This document states clearly that when a DNS resolver receives a response with response code NXDOMAIN, it means that the name in the question section AND ALL THE NAMES UNDER IT do not exist.

REMOVE BEFORE PUBLICATION: this document should be discussed in the IETF DNSOP (DNS Operations) group, through its mailing list. The source of the document, as well as a list of open issues, is currently kept at Github [1].

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 June 1, 2016.

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

Internet-Draft

NXDOMAIN cut

November 2015

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
1.1.	Requirements Terminology	3
2.	Rules	3
3.	Benefits	4
4.	Possible issues	4
5.	Future work	5
6.	IANA Considerations	6
7.	Security Considerations	6
8.	Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION	6
9.	Acknowledgments	7
10.	References	7
10.1.	Normative References	7
10.2.	Informative References	8
	Authors' Addresses	8

[1.](#) Introduction and background

The DNS protocol [[RFC1035](#)] defines response code 3 as "Name Error", or "NXDOMAIN", i.e. the queried domain name does not exist in the DNS. Since domain names are represented as a tree of labels ([\[RFC1034\], Section 3.1](#)), non-existence of a node implies non-existence of the entire sub-tree rooted at this node.

The DNS iterative resolution algorithm precisely interprets the NXDOMAIN signal in this manner. If it encounters an NXDOMAIN response code from an authoritative server, it immediately stops iteration and returns the NXDOMAIN response to the querier.

However, in virtually all existing resolvers, a cached NXDOMAIN is not considered "proof" that there can be no child domains underneath. This is due to an ambiguity in [[RFC1034](#)] that failed to distinguish ENT (empty nonterminal domain names, [[I-D.ietf-dnsop-dns-terminology](#)]) from nonexistent names. For DNSSEC, the IETF had to distinguish this case ([\[RFC4035\], section 3.1.3.2](#)), but the implication on non-DNSSEC resolvers wasn't fully realized.

This document specifies that an NXDOMAIN response for a domain name means that no child domains underneath the queried name exist either. And furthermore, that DNS resolvers should interpret cached NXDOMAIN responses in this manner. Since the domain names are organized in a

tree, it is a simple consequence of the tree structure: non-existence of a node implies non-existence of the entire sub-tree rooted at this node.

[1.1.](#) Requirements Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

[2.](#) Rules

When searching downward in its cache, an iterative caching DNS resolver SHOULD stop searching if it encounters a cached NXDOMAIN. The response to the triggering query should be NXDOMAIN.

TODO The next paragraph is challenged because too implementation-oriented. Should we just keep the first paragraph? My concern is that some resolvers may have an implementation of the cache made of a tree plus a hashed index and therefore won't "search downward" if they have a cached answer.

When an iterative caching DNS resolver stores an NXDOMAIN in its cache, all names and RRsets at or below that node SHOULD be deleted since they will have become unreachable. "Deleted" means that subsequent requests for these names will yield NXDOMAIN. [TODO: currently under discussion, some people find it dangerous. MAY instead of SHOULD? Only if the NXDOMAIN is DNSSEC-validated? Perhaps the resolver could be configured to not "bulk delete" TLDs (or root). See [Section 7.](#)] [TODO: currently under discussion, some people find it costly for the resolver. A large purge also causes the resolver to do a lot of (possibly CPU intensive) work, and could also affect the traffic levels between a recursive and authoritatives. Perhaps a cache may want to limit the number of nodes/records that it deletes per NXDOMAIN response.]

By implication, a stream of queries foo.example, then bar.foo.example, where foo.example does not exist would normally cause both queries to be forwarded to example's nameservers. Following this recommended practice of "NXDOMAIN cut", the second query and indeed any other query for names at or below foo.example would not be forwarded.

These rules replace the second paragraph of [section 5 of \[RFC2308\]](#). Otherwise, [\[RFC2308\]](#) applies unchanged, and the fact that a subtree does not exist is not forever: the NXDOMAIN is cached only for the duration of the "negative TTL" ([section 3](#) of xref target="[RFC2308](#)"/>).

Validating resolvers need to select the necessary NSEC or NSEC3 records and include them in the AUTHORITY section of the response to provide authenticated denial of existence for names underneath the NXDOMAIN boundary.

Warning: because of [\[RFC6604\]](#), the name whose existence is denied by the NXDOMAIN is not always the QNAME. If there is a chain of CNAME (or DNAME), the name which does not exist is the last of the chain. TODO: find a dedicated terminology such as "NXDOMAINed name" or "denied domain" instead of "QNAME".

[3.](#) Benefits

The main benefit is a better efficiency of the caches. In the example above, we send only one query instead of two, the second one being answered from the cache.

The correct behavior (in [\[RFC1034\]](#) and made clearer in this document) is specially useful when combined with QNAME minimisation [\[I-D.ietf-dnsop-qname-minimisation\]](#) since it will allow to stop searching as soon as a NXDOMAIN is encountered.

NXDOMAIN cut may also help mitigate certain types of random QNAME attacks [\[joost-dnsterior\]](#) [\[balakrichenan-dafa888\]](#), where there is a fixed suffix which does not exist. In these attacks against the authoritative name server, queries are sent to resolvers for a QNAME composed of a fixed suffix ("dafa888.wf" in one of the articles above), which is typically nonexistent, and a random prefix, different for each request. A resolver receiving these requests have

to forward them to the authoritative servers. With NXDOMAIN cut, we would just have to send to the resolver a query for the fixed suffix, the resolver would get a NXDOMAIN and then would stop forwarding the queries. (It would be better if the SOA record in the NXDOMAIN response were sufficient to find the non-existing domain but this is more delicate, see [Section 5](#).)

Since the principles set in this document are so great, why are the rules of [Section 2](#) SHOULD and not MUST? This is because some resolver may have a cache which is NOT organized as a tree (but, for instance, as a dictionary) and therefore have a good reason to ignore this.

[4](#). Possible issues

Let's assume the TLD example exists but foobar.example is not delegated (so the example's name servers will reply NXDOMAIN for a query about anything.foobar.example). A system administrator decides to name the internal machines of his organization under

office.foobar.example and use a trick of his resolver to forward requests about this zone to his local authoritative name servers. NXDOMAIN cut would create problems here, since, depending on the order of requests to the resolver, it may have cached the NXDOMAIN from example and therefore "deleted" everything under. This document assumes that such setup is rare and does not need to be supported.

Another issue that may happen: today, we see broken authoritative name servers which reply to ENT ([\[I-D.ietf-dnsop-dns-terminology\]](#), section 6) with NXDOMAIN instead of the normal NODATA ([\[I-D.ietf-dnsop-dns-terminology\]](#), section 3).

RFC-EDITOR: REMOVE THE PARAGRAPH BEFORE PUBLICATION. An example today is mta2._domainkey.cbs.nl (which exists) where querying _domainkey.cbs.nl yields NXDOMAIN. Another example is www.upenn.edu, redirected to www.upenn.edu-dscg.edgesuite.net while a query for edu-dscg.edgesuite.net returns NXDOMAIN.

Such name servers are definitely broken and have always been. They MUST be fixed. Given the advantages of NXDOMAIN cuts, there is little reason to support this behavior.

5. Future work

TODO: drop this section entirely? Or just downgrade it to an appendix "why can't we just use the owner name of the returned SOA"?

In this document, we deduce the non-existence of a domain only for NXDOMAIN answers where the QNAME was this exact domain. If a resolver sends a query to the name servers of the TLD example, and asks the MX record for www.foobar.example, and receives a NXDOMAIN, it can only register the fact that www.foobar.example (and everything underneath) does not exist. Even if the accompanying SOA record is for example only, one cannot infer that foobar.example is nonexistent. The accompanying SOA indicates the apex of the zone, not the closest existing domain name.

RFC-EDITOR: REMOVE BEFORE PUBLICATION: to use a real example today, ask the authoritative name servers of the TLD fr about anything.which.does.not.exist.gouv.fr. The SOA will indicate fr (the apex) even while gouv.fr does exist (there is no zone cut between gouv.fr and fr).

In the future, deducing the non-existence of a node from the SOA in the NXDOMAIN reply may certainly help with random qnames attacks but this is out-of-scope for this document. It would require to address the problems mentioned in the previous paragraph. A possible solution would be, when receiving a NXDOMAIN with a SOA which is more

than one label up in the tree, to send requests for the domains which are between the QNAME and the owner name of the SOA. (A resolver which does DNSSEC validation or QNAME minimisation will need to do it, anyway.)

TODO a mention of [[I-D.fujiwara-dnsop-nsec-aggressiveuse](#)]? Unlike NXDOMAIN cut, it requires DNSSEC but it is more powerful since it can synthesize NXDOMAINs.

6. IANA Considerations

This document has no actions for IANA.

7. Security Considerations

The technique described here may help against a denial-of-service attack named "random qnames" and described in [Section 3](#). Apart from that, it is believed to have no security consequences.

If a resolver does not validate the answers with DNSSEC, it can of course be poisoned with a false NXDOMAIN, thus "deleting" a part of the domain name tree. This denial-of-service attack is already possible with the rules of this document (but "NXDOMAIN cut" may increase its effects). The only solution is to use DNSSEC.

[8](#). Implementation status - RFC EDITOR: REMOVE BEFORE PUBLICATION

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [[RFC6982](#)]. The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalog of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [[RFC6982](#)], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

As of today, all existing DNS resolvers are conservative: they consider a NXDOMAIN as only significant for the name itself, not for the names under. All current recursive servers will upstream a query for out-of-cache sub.example.com even if their cache contains an NXDOMAIN for example.com.

[9](#). Acknowledgments

The text of this document was mostly copied from [[I-D.vixie-dnsexp-resimprove](#)], section 3. Thanks to its authors, Paul Vixie, Rodney Joffe and Frederico Neves.

Thanks to Duane Wessels, Tony Finch and Jinmei Tatuya for fact checking and explanations.

[10.](#) References

[10.1.](#) Normative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), DOI 10.17487/RFC1034, November 1987, <<http://www.rfc-editor.org/info/rfc1034>>.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), DOI 10.17487/RFC1035, November 1987, <<http://www.rfc-editor.org/info/rfc1035>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/[RFC2119](#), March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](#), DOI 10.17487/RFC2308, March 1998, <<http://www.rfc-editor.org/info/rfc2308>>.
- [RFC6604] Eastlake 3rd, D., "xNAME RCODE and Status Bits Clarification", [RFC 6604](#), DOI 10.17487/RFC6604, April 2012, <<http://www.rfc-editor.org/info/rfc6604>>.
- [I-D.ietf-dnsop-dns-terminology] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", [draft-ietf-dnsop-dns-terminology-05](#) (work in progress), September 2015.

[10.2.](#) Informative References

- [RFC4035] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "Protocol Modifications for the DNS Security Extensions", [RFC 4035](#), DOI 10.17487/RFC4035, March 2005, <<http://www.rfc-editor.org/info/rfc4035>>.
- [RFC6982] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", [RFC 6982](#), DOI 10.17487/RFC6982, July 2013, <<http://www.rfc-editor.org/info/rfc6982>>.
- [I-D.vixie-dnsexst-resimprove]
Vixie, P., Joffe, R., and F. Neves, "Improvements to DNS Resolvers for Resiliency, Robustness, and Responsiveness", [draft-vixie-dnsexst-resimprove-00](#) (work in progress), June 2010.
- [I-D.ietf-dnsop-qname-minimisation]
Bortzmeyer, S., "DNS query name minimisation to improve privacy", [draft-ietf-dnsop-qname-minimisation-07](#) (work in progress), October 2015.
- [I-D.fujiwara-dnsop-nsec-aggressiveuse]
Fujiwara, K. and A. Kato, "Aggressive use of NSEC/NSEC3", [draft-fujiwara-dnsop-nsec-aggressiveuse-02](#) (work in progress), October 2015.
- [joost-dnsterror]
Joost, M., "About DNS Attacks and ICMP Destination Unreachable Reports", December 2014, <<http://www.michael-joost.de/dnsterror.html>>.
- [balakrichenan-dafa888]
Balakrichenan, S., "Disturbance in the DNS - "Random qnames", the dafa888 DoS attack"", October 2014, <<https://indico.dns-oarc.net/event/20/session/3/contribution/37>>.

Authors' Addresses

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/>

Shumon Huque
Verisign labs
12061 Bluemont Way
Reston 20190
USA

Email: shuque@verisign.com
URI: <http://www.verisignlabs.com/>

