DNSOP
Internet-Draft
Intended status: Best Current Practice
Expires: June 1, 2018

A. Sullivan
Oracle
J. Abley
Snake Hill Labs

November 28, 2017

Please See Below: Use Only Downward Referrals in the DNS
draft-sullivan-dnsop-refer-down-00

Abstract

A server in the Domain Name System can use a mechanism called "referral" to indicate that the server is not authoritative for a given zone, and to redirect the query to another, more appropriate server. The mechanism was originally specified such that a referral might be to any location in the DNS. Operational experience indicates dubious value to referrals other than those to zones below the zones for which a server is authoritative. This memo therefore recommends such referrals and discourages other kinds of referrals.

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 https://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, 2018.

Table of Contents

1. Introduction 2
1.1. Terminology 3
2. Referrals 3
2.1. Downward Referrals 3
2.2. Upward Referrals 4
2.3. Negative Consequences of Upward Referrals 5
2.4. Alternatives to Upward Referrals 6
2.4.1. NODATA 6
2.4.2. SERVFAIL 6
2.4.3. NXDOMAIN 6
2.4.4. REFUSED 6
2.5. Recommendations 7
3. Acknowledgements 7
4. IANA Considerations 7
5. References 7
5.1. Normative References 7
5.2. Informative References 8
Appendix A. Discussion Venue 8
Appendix B. Change History 9
Authors' Addresses 9

1. Introduction

The Domain Name System (DNS) divides parts of the domain name space "into units called 'zones'" ([RFC1034], Section 2.4). The answers for data in these zones are (ultimately) provided by authoritative servers. In the Internet context, for any given query, there is a set of authoritative servers that can provide an authoritative answer in response to that query.

Sometimes, however, a server receives a query for which it is not authoritative. If such a server does not offer recursion, the server might return a response that refers to another set of servers on the Internet. This response is called a "referral".

There are two categories of referral response. One of them indicates a delegation in the DNS, and is a basic part of how the DNS functions. Without such delegation responses, the distributed nature of the DNS is impossible. They may be thought of as "downward" referrals because they refer to a zone somewhere beneath the zone for which the server is authoritative. Other referrals are for zones

where the server is neither authoritative for the zone of the QNAME, nor for any zone that might be an ancestor of the zone containing the QNAME. These referrals might be thought of as "off-tree" referrals, because the server is not authoritative for any part of the tree containing the QNAME.

Historically, authoritative servers that received an off-tree query would reply with an "upward referral", usually to the root zone; these were sometimes called a "root referral". Such referrals have turned out to be undesirable in practice. This memo recommends that servers not provide upward referrals, and instead should respond to such queries in some other way.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Unfamiliar DNS-related terms are likely to be found in RFC 7719 [RFC7719], and the reader is assumed to be familiar with that vocabulary.

2. Referrals

Referrals are defined as part of the algorithm for a name server ([RFC1034], section 4.3.2, henceforth "the algorithm"). Referrals only happen when the RD bit is clear in the query or the server does not offer recursion (or both). There are different possible interpretations of the algorithm; one's interpretation will affect which kinds of referral one thinks acceptable.

A referral contains an empty answer section. It contains the NS RRset for the referred-to zone in the authority section. It may contain RRs that provide addresses in the additional section. The AA bit is clear.

2.1. Downward Referrals

The first kind of referral is downward, and is uncontroversial. Step 2 of the matching algorithm evaluates whether the name server is authoritative for some zone that is an ancestor for the QNAME. (If the QNAME exactly matches, then that zone is the "ancestor". This is a slightly awkward usage of "ancestor", but makes sense due to the distinction between a zone and the matching owner name inside the

zone.) If there is such a zone, then the algorithm moves to step 3; otherwise, it moves to step 4.

In step 3, the server matches label by label in the zone until matching terminates. Step 3(b) of the matching algorithm says, "If a match would take us out of the authoritative data, we have a referral. This happens when we encounter a node with NS RRs marking cuts along the bottom of a zone." Such a referral is called "downward" because the referral is of necessity to a part of the namespace beneath the zone for which the server is generating a response. In other words, if the server is authoritative for the zone example.com, the referral needs to be to the NS records of some subordinate zone in the domain name space.

Downward referrals are necessary for the DNS to function. They are the mechanism by which delegation happens.

2.2. Upward Referrals

The second kind of referral is often called an "upward" referral, because it is often a referral to the name servers for the root zone (perforce above everything else in the domain name space), though in principle the referral could be elsewhere in the domain name space. Step 4 of the algorithm says, "If there was no delegation from authoritative data, look for the best one from the cache, and put it in the authority section." Returning this kind of referral under normal operational conditions is somewhat more controversial than a downward referral, because it is not clear that it is necessary for the operation of the DNS.

There are only two cases where upward referrals are possible:

1. The server offers recursive service, and it cannot provide an authoritative answer or a downward referral, but the query was received with the RD bit clear.

2. The server does not offer recursive service, and it cannot provide either an answer or a downward referral in response to the query.

The first of these is plainly required by step 4 of the algorithm, and should therefore be uncontroversial. In normal operation, however, this case appears to be unusual. A resolver that was using such a server for full-service DNS resolution would normally query with the RD bit set. A resolver that did not expect recursion would likely only send a QNAME for which the server could provide an authoritative answer or a downward referral; it is unclear why the query would be sent to the server at all otherwise. Such queries are

known to occur sometimes, for example when troubleshooting, but they do not appear to be normal according to the protocol.

The second case is controversial because the server, which only provides authoritative answers, must somehow have some data in a cache in order to return anything in the authority section. The controversy arises because of the question of whether the server ought to have such data. This amounts to a question of whether a server that only provides authoritative answers should ever have a cache.

On the one hand, it would seem that such a server should not have a cache, because it does not have a resolver side that populates such a cache. Moreover, the SBELT structure (see [RFC1034], section 5.3.2) is defined only for resolvers and not for servers. So a server that only provides authoritative answers has no reason even to have configured in the SBELT structure a list of servers from which to start (in resolvers, this is often the "root hints" file). On the other hand, there is no requirement that a given name server should not provide both authoritative service and recursive service. Moreover, even a server that provides no recursive service to others may need to perform resolution for its own purposes, and therefore might have need of the SBELT structure. So, depending on one's reading of the algorithm, either upward referrals should not be returned from such a server and are a sign of misconfiguration, or else they will be a normal part of operation.

Upward referrals, and particularly root referrals, were once regarded as a useful mechanism to indicate lame delegation [RFC1912]. That use turned out to create some difficulties (see Section 2.3, below).

2.3. Negative Consequences of Upward Referrals

Upward referrals have some negative consequences. The most obvious of them is that they are not in-domain records, and therefore they should not be accepted in any case according to RFC 5452 [RFC5452], section 6. This means that an upward referral response is just extra traffic, because the querying resolver will need to find those records from an authoritative source anyway. Moreover, upward referral response messages can be considerably larger than the query message that causes them, making them a useful amplifier when used in reflector attacks [RFC5358].

Upward referrals can be part of a referral loop, and the algorithm does not specify how or when to terminate such a loop. The use of upward referrals to indicate lame delegations exhibits this weakness.

2.4. Alternatives to Upward Referrals

It is possible for a server to send some other response than an upward referral, when an upward referral might have been generated under the algorithm. There are several alternatives, each of which has advantages and disadvantages.

2.4.1. NODATA

A name server that had no information at all in a cache (including the SBELT structure) would complete step 4 of the algorithm having added nothing to the authority section in the response. It would exit step 6 of the algorithm having created an empty response (except for the query that was copied from the original query message). This is a type 3 NODATA response [RFC2308]. A disadvantage of returning such a message is that it is unlikely to cause the query source to stop querying the nameserver for that name, because type 3 NODATA responses are not cached (see [RFC2308], section 5).

2.4.2. SERVFAIL

RCODE 2, Server Failure, indicates that a server cannot process the query due to a problem with the name server. Some operators adopt the position that the name server would normally provide an upward referral, except that it has been configured not to. Therefore, the server can return RCODE 2. Others argue, however, that there is nothing wrong with the server; and that, moreover, the use of RCODE 2 in DNSSEC (see [RFC4035]) means that this RCODE is already overloaded enough. Some interpretations of RCODE 2 by resolvers invites subsequent retries to the same server, which may not always be desirable.

2.4.3. NXDOMAIN

RCODE 3, Name Error or NXDOMAIN, indicates that the domain name does not exist. Some operators use RCODE 3 instead of producing upward referrals. But since RCODE 3 is supposed to be "[m]eaningful only for responses from an authoritative name server" ([RFC1035] section 4.1.1) and since by definition the upward referral can only happen in a case where the name server is not authoritative, this use appears to be inconsistent with the protocol.

2.4.4. REFUSED

RCODE 5, Refused, indicates that the server "refuses to perform the specified operatio for policy reasons." ([RFC1035], section 4.1.1) Some operators adopt a policy of refusing to perform upward referrals, and so return RCODE 5 to queries that would otherwise

cause such referrals. There are some resolvers, however, that interpret RCODE 5 to mean that the resolver itself, rather than the query sent, is what causes the Refused response. Those resolvers will not attempt to query the server again (or not for some period of time), running the risk of outages in domains for which the server is authoritative and would provide a response.

2.5. Recommendations

A name server that only provides authoritative service SHOULD NOT return upward referrals under any circumstances. Such a name server SHOULD provide either RCODE 2 or RCODE 5 in response. A name server MUST NOT return RCODE 3 except for names for which it can provide authoritative answer that the name does not exist.

A name server that provides recursive service MAY provide upward referrals when replying to a query with the RD bit clear, or it MAY refuse to provide upward referrals just as though it provided only authoritative service. Operators should note that upward referrals might provide a modest troubleshooting advantage for recursive servers, but this should be weighed against the advantages of removing upward referrals as one of the available tools of attackers on Internet infrastructure.

3. Acknowledgements

This memo has benefitted from the comments of Stephane Bortzmeyer, Robert Edmonds, Tony Finch, Evan Hunt, John Kristoff, Dave Lawrence, Edward Lewis, Matthew Pounsett, and Paul Vixie.

4. IANA Considerations

This memo makes no requests of IANA.

[[CREF1: Note in draft: this section can be removed by the RFC Editor if the document is ever published as an RFC.]]

5. References

5.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

5.2. Informative References

[RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC1912] Barr, D., "Common DNS Operational and Configuration Errors", RFC 1912, DOI 10.17487/RFC1912, February 1996, <https://www.rfc-editor.org/info/rfc1912>.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, DOI 10.17487/RFC2308, March 1998, <https://www.rfc-editor.org/info/rfc2308>.

[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, <https://www.rfc-editor.org/info/rfc4035>.

[RFC5358] Damas, J. and F. Neves, "Preventing Use of Recursive Nameservers in Reflector Attacks", BCP 140, RFC 5358, DOI 10.17487/RFC5358, October 2008, <https://www.rfc-editor.org/info/rfc5358>.

[RFC5452] Hubert, A. and R. van Mook, "Measures for Making DNS More Resilient against Forged Answers", RFC 5452, DOI 10.17487/RFC5452, January 2009, <https://www.rfc-editor.org/info/rfc5452>.

[RFC7719] Hoffman, P., Sullivan, A., and K. Fujiwara, "DNS Terminology", RFC 7719, DOI 10.17487/RFC7719, December 2015, <https://www.rfc-editor.org/info/rfc7719>.

Appendix A. Discussion Venue

This Internet-Draft is discussed on the DNS Operations Working Group list: dnsop@ietf.org.

Appendix B. Change History

Note to RFC Editor: this section should be removed prior to publication as an RFC.

00:

* Initial version

Authors' Addresses

Andrew Sullivan
Oracle

Email: andrew.s.sullivan@oracle.com

Joe Abley
Snake Hill Labs
300-184 York Street
London, ON N6A 1B5
Canada

Email: jabley@shl.io