Network Working Group
Internet-Draft
Updates: 7050 (if approved)
Intended status: Standards Track
Expires: May 3, 2017

S. Cheshire
D. Schinazi
Apple Inc.

October 30, 2016

Special Use Domain Name 'ipv4only.arpa' draft-cheshire-sudn-ipv4only-dot-arpa-04

Abstract

The specification for how a client discovers its network's NAT64 prefix [RFC7050] defines the special name 'ipv4only.arpa' for this purpose, but declares it to be a non-special name in that specification's Domain Name Reservation Considerations section.

Consequently, despite the well articulated special purpose of the name, as of July 2016 'ipv4only.arpa' still does not appear as one of the names with special properties that are recorded in the Special-Use Domain Names registry.

This document formally declares the actual special properties of the name, and adds similar declarations for the corresponding reverse mapping names.

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 May 3, 2017.

Copyright Notice

1. Introduction

The specification for how a client discovers its network's NAT64 prefix [RFC7050] defines the special name 'ipv4only.arpa' for this purpose, but declares it to be a non-special name in that specification's Domain Name Reservation Considerations section.

Consequently, despite the well articulated special purpose of the name, as of July 2016 'ipv4only.arpa' still does not appear as one of the names with special properties that are recorded in the Special-Use Domain Names registry [SUDN].

This document formally declares the actual special properties of the name. This document also adds similar declarations for the corresponding reverse mapping names.

2. Specialness of 'ipv4only.arpa'

The hostname 'ipv4only.arpa' is peculiar in that it was never intended to be treated like a normal hostname.

A typical client never looks up the IPv4 address records for 'ipv4only.arpa', because it is already known, by specification [RFC7050], to have exactly two IPv4 address records, 192.0.0.170 and 192.0.0.171. No client ever has to look the name in order to learn those two addresses.

In contrast, clients often look up the IPv6 AAAA address records for 'ipv4only.arpa', which is contrary to general DNS expectations, given that it is already known, by specification [RFC7050], that no such IPv6 AAAA address records exist. And yet, clients expect to receive, and do in fact receive, positive answers for these IPv6 AAAA address records that are known to not exist.

This is clearly not a typical DNS name. In normal operation, clients never query for the two records that do in fact exist; instead they query for records that are known to not exist, and then get positive answers to those abnormal queries. Clients are using DNS to perform queries for this name, but they are certainly not using DNS to learn legitimate answers from the name's legitimate authoritative server. Instead, these clients have, in effect, co-opted the DNS protocol as an impromptu client-to-middlebox communication protocol, to communicate with the NAT64/DNS64 [RFC6146][RFC6147] gateway, if present, and request that it disclose the prefix it is using for IPv6 address synthesis.

It is this use of specially-crafted DNS queries as an impromptu client-to-middlebox communication protocol that makes the name 'ipv4only.arpa' most definitely a special name, and one that should, in the spirit of openness and honesty, be listed in IANA's registry along with other DNS names that have special uses [SUDN].

3. Consequences of 'ipv4only.arpa' previously being declared unspecial

As a result of the original specification [RFC7050] not formally declaring 'ipv4only.arpa' to have special properties, there was no mandate for any server software to treat this name specially. Consequently, queries for this name had to be handled normally, resulting in unnecessary queries to the authoritative 'arpa' name servers.

Having millions of devices around the world issue these queries generated pointless additional load on the authoritative 'arpa' name servers, which was completely unnecessary when the name 'ipv4only.arpa' is defined, by Internet Standard, to have exactly two IPv4 address records, 192.0.0.170 and 192.0.0.171, and no other records of any type.

Also, at times, for reasons that are as yet unclear, the authoritative 'arpa' name servers have been observed to be slow or unresponsive. The failures of these 'ipv4only.arpa' queries result in unnecessary failures of software that depends on them for DNS64 [RFC6147] address synthesis.

Even when the authoritative 'arpa' name servers are operating correctly, having to perform an unnecessary query to obtain an answer that is already known in advance can add precious milliseconds of delay for no reason.

This document leverages this operational experience to update the Domain Name Reservation Considerations section [RFC6761] of the earlier specification [RFC7050] with one that accurately lists the actual special properties of the name 'ipv4only.arpa' so that software can legitimately make appropriate performance and reliability optimizations.

4. Security Considerations

Hard-coding the known answers for 'ipv4only.arpa' queries in recursive/caching DNS servers reduces the risk of malicious devices intercepting those queries and returning incorrect answers, particularly in the case of recursive/caching DNS servers that do not perform DNSSEC validation.

One of the known concerns with DNS64 [RFC6147] is that it interferes with DNSSEC. DNSSEC may cryptographically assert that a name has no IPv6 AAAA records, while at the same time DNS64 address synthesis is contradicting this and claiming that IPv6 AAAA records do exist.

Section 3 of the DNS64 specification [RFC6147] discusses this:

> ... DNS64 receives a query with the DO bit set and
> the CD bit set. In this case, the DNS64 is supposed
> to pass on all the data it gets to the query initiator.
> This case will not work with DNS64, unless the
> validating resolver is prepared to do DNS64 itself.

The NAT64 Prefix Discovery specification [RFC7050] provides the mechanism for the query initiator to learn the NAT64 prefix so that it can do its own validation and DNS64 synthesis as described above. With this mechanism the client can (i) interrogate the local NAT64/DNS64 gateway with an 'ipv4only.arpa' query to learn the IPv6 address synthesis prefix, (ii) query for the (signed) IPv4 address records itself, and then (iii) perform its own IPv6 address synthesis locally, combining the IPv6 address synthesis prefix learned from the local NAT64/DNS64 gateway with the secure DNSSEC-signed data learned from the global Domain Name System.

It is conceivable that over time, if DNSSEC is successful, the majority of clients could move to this validate-and-synthesize-locally model, which reduces the DNS64 machinery to the vestigial role of simply responding to the 'ipv4only.arpa' query to report the local IPv6 address synthesis prefix. In no case does the client care what answer(s) the authoritative 'arpa' name servers might give for that query. The 'ipv4only.arpa' query is being used purely as a local client-to-middlebox communication message.

This approach is even more attractive if it doesn't create an additional dependency on the authoritative 'arpa' name servers to answer a query that is unnecessary because the NAT64/DNS64 gateway already knows the answer before it even issues the query. Avoiding this unnecessary query improves performance and reliability for the client, and reduces unnecessary load for the authoritative 'arpa' name servers.

5. IANA Considerations

[Once published, this should say]

IANA has recorded the following names in the Special-Use Domain Names registry [SUDN]:

ipv4only.arpa.
170.0.0.192.in-addr.arpa.
171.0.0.192.in-addr.arpa.

IANA has recorded the following IPv4 addresses in the IPv4 Special-Purpose Address Registry [SUv4]:

192.0.0.170
192.0.0.171

6. Domain Name Reservation Considerations

6.1. Conventions and Terminology Used in this Section

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this section are to be interpreted as described in "Key words for use in RFCs to Indicate Requirement Levels" [RFC2119].

6.2. ipv4only.arpa

The name 'ipv4only.arpa' is defined, by Internet Standard, to have two IPv4 address records with rdata 192.0.0.170 and 192.0.0.171.

When queried via a DNS64 [RFC6147] recursive/caching server, the name 'ipv4only.arpa' is also defined to have IPv6 AAAA records, with rdata synthesized from a combination of the NAT64 IPv6 prefix(es), and the IPv4 addresses 192.0.0.170 and 192.0.0.171. This can return more than one pair of IPv6 addresses if there are multiple NAT64 prefixes.

The name 'ipv4only.arpa' has no other DNS records of any type.

The name 'ipv4only.arpa' is special only to
(a) client software wishing to perform DNS64 address synthesis, and
(b) the DNS64 recursive/caching server responding to such requests.
These two considerations are listed in items 2 and 4 below:

1. Normal users should never have reason to encounter the 'ipv4only.arpa' domain name. If they do, they should expect queries for 'ipv4only.arpa' to result in the answers required by the specification [RFC7050]. Normal users have no need to know that 'ipv4only.arpa' is special.

2. Application software may explicitly use the name 'ipv4only.arpa' for NAT64/DNS64 address synthesis, and expect to get the answers required by the specification [RFC7050]. If application software encounters the name 'ipv4only.arpa' in the normal course of handling user input, the application software should resolve that name as usual and need not treat it in any special way.

3. Name resolution APIs and libraries SHOULD NOT recognize 'ipv4only.arpa' as special and SHOULD NOT treat it differently. Name resolution APIs SHOULD send queries for this name to their configured recursive/caching DNS server(s).

4. Recursive/caching DNS servers SHOULD recognize 'ipv4only.arpa' as special and SHOULD NOT, by default, attempt to look up NS records for it, or otherwise query authoritative DNS servers in an attempt to resolve this name. Instead, recursive/caching DNS servers SHOULD, by default, act as authoritative and generate immediate responses for all such queries.

   Traditional recursive/caching DNS servers that act as authoritative for this name MUST generate only the 192.0.0.170 and 192.0.0.171 responses for IPv4 address queries (DNS qtype "A"), and a negative ("no error no answer") response for all other query types.

All DNS64 recursive/caching DNS servers MUST generate the 192.0.0.170 and 192.0.0.171 responses for IPv4 address queries (DNS qtype "A"), the appropriate synthesized IPv6 address record responses for IPv6 address queries (DNS qtype "AAAA"), and a negative ("no error no answer") response for all other query types.

This local self-contained generation of these responses is to avoid placing unnecessary load on the authoritative 'arpa' name servers.

Example configurations for BIND 9 showing how to achieve these results are given in Appendix A.

5. Traditional authoritative DNS server software need not recognize 'ipv4only.arpa' as special or handle it in any special way. Recursive/caching DNS servers SHOULD routinely act as authoritative for this name and return the results described above. Only the administrators of the 'arpa' namespace need to explicitly configure their actual authoritative name servers to be authoritative for this name and to generate the appropriate answers; all other authoritative name servers will not be configured to know anything about this name and will reject queries for it, as they would reject queries for any other name about which they have no information.

6. Generally speaking, operators of authoritative DNS servers need not know anything about the name 'ipv4only.arpa', just as they don't need to know anything about any other names they are not responsible for. Operators of authoritative DNS servers who are configuring their name servers to be authoritative for this name MUST understand that 'ipv4only.arpa' is a special name, with records rigidly specified by Internet Standard (generally this applies only to the administrators of the 'arpa' namespace).

7. DNS Registries/Registrars need not know anything about the name 'ipv4only.arpa', just as they don't need to know anything about any other name they are not responsible for. Only the administrators of the 'arpa' namespace need to be aware of this name's purpose and how it should be configured.

6.3. 170.0.0.192.in-addr.arpa and 171.0.0.192.in-addr.arpa

Since the IPv4 addresses 192.0.0.170 and 192.0.0.171 are defined to be special, and are listed in the IPv4 Special-Purpose Address Registry [SUv4], the corresponding reverse mapping names in the in-addr.arpa domain are similarly special.

The name '170.0.0.192.in-addr.arpa' is defined, by Internet Standard, to have only a single DNS record, type PTR, with rdata 'ipv4only.arpa'.

The name '171.0.0.192.in-addr.arpa' is defined, by Internet Standard, to have only a single DNS record, type PTR, with rdata 'ipv4only.arpa'.

Practically speaking these two names are rarely used, but to the extent that they may be, they are special only to recursive/caching DNS servers as described in item 4 below:

1. Normal users should never have reason to encounter these two reverse mapping names. However, if they do, queries for these reverse mapping names should return the expected answer 'ipv4only.arpa'. Normal users have no need to know that these reverse mapping names are special.

2. Application software SHOULD NOT recognize these two reverse mapping names as special, and SHOULD NOT treat them differently. For example, if the user were to issue the Unix command "host 192.0.0.170" then the "host" command should issue the query as usual and display the result that is returned.

3. Name resolution APIs and libraries SHOULD NOT recognize these two reverse mapping names as special and SHOULD NOT treat them differently. Name resolution APIs SHOULD send queries for these names to their configured recursive/caching DNS server(s).

4. Recursive/caching DNS servers SHOULD recognize these two reverse mapping names as special and SHOULD NOT, by default, attempt to look up NS records for them, or otherwise query authoritative DNS servers in an attempt to resolve them. Instead, recursive/caching DNS servers SHOULD, by default, act as authoritative and generate immediate responses for all such queries.

   Recursive/caching DNS servers that act as authoritative for these names MUST generate only the 'ipv4only.arpa' response for PTR queries, and a negative ("no error no answer") response for all other query types. This local self-contained generation of these responses is to avoid placing unnecessary load on the

authoritative 'in-addr.arpa' name servers.

5. Traditional authoritative DNS server software need not recognize these two reverse mapping names as special or handle them in any special way.
   As a practical matter, only the administrators of the '192.in-addr.arpa' namespace will configure their name servers to be authoritative for these names and to generate the appropriate answers; all other authoritative name servers will not be configured to know anything about these names and will reject queries for them as they would reject queries for any other name about which they have no information.

6. Generally speaking, operators of authoritative DNS servers need not know anything about these two reverse mapping names, just as they don't need to know anything about any other names they are not responsible for. Operators of authoritative DNS servers who are configuring their name servers to be authoritative for this name MUST understand that these two reverse mapping names are special, with answers specified by Internet Standard (generally this applies only to the administrators of the '192.in-addr.arpa' namespace).

7. DNS Registries/Registrars need not know anything about these two reverse mapping names, just as they don't need to know anything about any other name they are not responsible for. Only the administrators of the '192.in-addr.arpa' namespace need to be aware of the purpose of these two names.

6.4. ip6.arpa Reverse Mapping PTR Records

For all IPv6 addresses synthesized by the NAT64 gateway, the DNS64 recursive/caching server is responsible for synthesizing the appropriate ip6.arpa reverse mapping PTR records, if it chooses to do so. The same applies to the synthesized IPv6 addresses corresponding to the IPv4 addresses 192.0.0.170 and 192.0.0.171.

Generally a DNS64 recursive/caching server synthesizes appropriate ip6.arpa reverse mapping PTR records by extracting the embedded IPv4 address from the encoded IPv6 address, performing a reverse mapping query for that IPv4 address, and then synthesizing a corresponding ip6.arpa reverse mapping PTR record containing the same rdata.

In the case of synthesized IPv6 addresses corresponding to the IPv4 addresses 192.0.0.170 and 192.0.0.171, the DNS64 recursive/caching server does not issue mapping queries for those IPv4 addresses, but instead, according to rule 3 above, immediately returns the answer 'ipv4only.arpa'.

7. References

7.1. 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, <http://www.rfc-editor.org/info/rfc2119>.

[RFC6146] Bagnulo, M., Matthews, P., and I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", RFC 6146, DOI 10.17487/RFC6146, April 2011, <http://www.rfc-editor.org/info/rfc6146>.

[RFC6147] Bagnulo, M., Sullivan, A., Matthews, P., and I. van Beijnum, "DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", RFC 6147, DOI 10.17487/RFC6147, April 2011, <http://www.rfc-editor.org/info/rfc6147>.

[RFC6761] Cheshire, S. and M. Krochmal, "Special-Use Domain Names", RFC 6761, DOI 10.17487/RFC6761, February 2013, <http://www.rfc-editor.org/info/rfc6761>.

[RFC7050] Savolainen, T., Korhonen, J., and D. Wing, "Discovery of the IPv6 Prefix Used for IPv6 Address Synthesis", RFC 7050, DOI 10.17487/RFC7050, November 2013, <http://www.rfc-editor.org/info/rfc7050>.

7.2. Informative References

[SUDN] "Special-Use Domain Names Registry", <https://www.iana.org/assignments/special-use-domain-names/>.

[SUv4] "IANA IPv4 Special-Purpose Address Registry", <https://www.iana.org/assignments/iana-ipv4-special-registry/>.

Appendix A. Example BIND 9 Configuration

A BIND 9 recursive/caching DNS server could be configured to act as authoritative for the appropriate names as follows.

In /etc/named.conf the following lines are added:

```
zone "ipv4only.arpa"            { type master; file "ipv4only"; };
zone "170.0.0.192.in-addr.arpa" { type master; file "ipv4only-r"; };
zone "171.0.0.192.in-addr.arpa" { type master; file "ipv4only-r"; };
```

The file /var/named/ipv4only is created with the following content. (The lines marked "Only for DNS64 server" are omitted on a standard recursive/caching DNS server.)

```
$TTL 86400              ; Default TTL 24 hours
@ IN SOA nameserver.example. admin.nameserver.example. (
         2016052400     ; Serial
         7200           ; Refresh ( 7200 = 2 hours)
         3600           ; Retry   ( 3600 = 1 hour)
         15724800       ; Expire  (15724800 = 6 months)
         60             ; Minimum
         )
@ IN NS  nameserver.example.

@ IN A     192.0.0.170
@ IN A     192.0.0.171
@ IN AAAA  64:ff9b::192.0.0.170 ; Only for DNS64 server
@ IN AAAA  64:ff9b::192.0.0.171 ; Only for DNS64 server
```

The file /var/named/ipv4only-r is created with the following content:

```
$TTL 86400              ; Default TTL 24 hours
@ IN SOA nameserver.example. admin.nameserver.example. (
         2016052400     ; Serial
         7200           ; Refresh ( 7200 = 2 hours)
         3600           ; Retry   ( 3600 = 1 hour)
         15724800       ; Expire  (15724800 = 6 months)
         60             ; Minimum
         )
@ IN NS  nameserver.example.

@ IN PTR ipv4only.arpa.
```

Authors' Addresses

Stuart Cheshire
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
USA

Phone: +1 408 974 3207
Email: cheshire@apple.com

David Schinazi
Apple Inc.
1 Infinite Loop
Cupertino, California 95014
USA

Phone: +1 669 227 9921
Email: dschinazi@apple.com