

Internationalized PTR Resource Record (IPTR)

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>

Abstract

This draft attempts to address the problem of how an IP address should be properly mapped to a set of internationalized domain names(iDNs). It is currently unspecified how a PTR record can be used for this purpose. In addition, the syntax of the PTR resource record may be too restrictive for such a mapping in a more culturally meaningful context. This document suggests a new TYPE called IPTR using EDNS0 and a mechanism to combine language information with such a mapping.

1. Introduction

Reverse mapping is a very important and essential function in the DNS. In today's Domain Name System, PTR RRs are used to support address-to-domain mappings. However, a current PTR RR does not provide support for proper address-to-iDN mappings, without certain modifications.

Modifying the PTR structure will also affect the current reverse mapping architecture. This document describes a new RR TYPE named IPTR to provide address-to-iDN mappings and it also specifies that on receiving of a IPTR query a name server should respond with all the corresponding IPTR RRs in one response. This document also specifies that an IPTR RR SHOULD refer to one primary iDN per language only.

1.1 Terminology

The key words "MUST", "SHALL", "REQUIRED", "SHOULD", "RECOMMENDED", and "MAY" in this document are to be interpreted as described in [RFC 2119](#) [[RFC2119](#)].

1.2 Background and Designs

When Internationalized Domain Names come into wide use, an Internet host is likely to have domain names in different languages. In today's Internet, because of the design of the PTR record and implementation of most resolvers, IP address to domain names mapping is limited to "one IP one domain name", the primary domain name of the host. This is more restrictive in a world of iDNs, for choosing one name in one particular language as the primary could have cultural implications. The authors also believe that putting language information into address-to-name mappings will be benifitial to future applications.

The design purpose of the IPTR RR type is to provide a mechanism that can map an IP address to all of the corresponding iDNs per language.

CNAME MUST continue to work for IPTR as it works now for PTR records. An IPTR RR SHOULD be limited to one primary iDN per language.

The behavior of a resolver on the use of IPTR will be specified in a seperate draft or a later version of this draft.

1.3 Functional Description

DNS query and responses involving IPTR type MUST have the following properties:

- When the QTYPE is IPTR, the corresponding iDNs SHOULD be returned in one response.
- The characters in the label MUST be encoded using UTF-8 [[RFC2279](#)].

- The entire label MUST be encoded EDNS [[RFC2671](#)].

2. IPTR definition

The structure of an IPTR RR is somewhat like the MX RR. :) In addition to the IP address in the IN-ADDR.ARPA domain and the domain name field (similar to a PTR RR), a new field called LANGUAGE has been defined. A domain name in an IPTR RR MUST be encoded in UTF8. Below is an example of an IPTR RR:

```
1.2.3.4.IN-ADDR.ARPA.    IPTR  "language" "name-in-utf8"
```

[RFC1766] describes the ISO 639/ISO 3166 conventions. A language name is always written in lower case, while country codes are written in upper case. The "language" field in an IPTR RR MUST follow the conventions defined in [[RFC1766](#)].

For Example:

```
4.3.2.1.IN-ADDR.ARPA.    IPTR    "zh-cn"    "name-in-utf8"
4.3.2.1.IN-ADDR.ARPA.    IPTR    "zh-tw"    "name-in-utf8"
4.3.2.1.IN-ADDR.ARPA.    IPTR    "ja-jp"    "name-in-utf8"
4.3.2.1.IN-ADDR.ARPA.    IPTR    "ko-kr"    "name-in-utf8"
```

The notion of canonical names and aliases described in 3.6.2 [[RFC1034](#)] must be preserved for IPTR record types. An IPTR RR SHOULD be limited to one primary iDN per language, similar to the a PTR RR.

3. IPTR on IPv6

Mapping IPv6 to iDNs can be similarly supported. This document recommends to continue using the IP6.INT domain defined in [[RFC1886](#)] for IPTR mappings. For example, the lookup corresponding to the address 4321:0:1:2:3:4:567:89ab would be:

```
b.a.9.8.7.6.5.0.4.0.0.0.3.0.0.0.2.0.0.0.1.0.0.0.0.0.0.1.2.3.4.IP6.INT.
IPTR  "language" "name-in-utf8"
```

4. Packet format for IPTR

EDNS0[RFC2671] is REQUIRED to implement IPTR.


```

      0          1          2          3          4
bits 0 1 2 3 4 5 6 7 8 9 0 1...9 0...8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 ...
      +-+-+-+-+-+-+//+-+-+//+-+-+//+-+-+//+-+-+//+-+-+//+-+-+
      |0 1|    ELT    |  LANGUAGE    |      Size      | IDN label... |
      +-+-+-+-+-+-+//+-+-+//+-+-+//+-+-+//+-+-+//+-+-+//+-+-+

```

LANGUAGE: An argument for IPTR to define the kind of languages
used in the following IDN label. The size is 2 octets.

ELT: To be defined.

5. IPTR query/response

When the QTYPE is IPTR in a query, all of the corresponding IPTR RRs SHOULD be returned in one response. DNS messages are limited to 512 octets or less in size when sent over UDP. Therefore, if all the RRs cannot fit in one UDP packet, this draft describe two solutions. One is for recent environment and the other is for the near future.

5.1 Transport

Today, DNS queries and responses are carried in UDP datagrams or over TCP connections. [\[RFC1034\]](#) specifies, IPTR RRSet is RECOMMENDED to be returned in one response. The size of a DNS message could exceed 512 octets, when multiple RRs are present. Therefore, this draft makes the two following recommendations.

- "Use UDP first, if UDP is not large enough then change to TCP" is RECOMMENDED.

The server MUST send back the response with the TC bit set. Then the resolver SHOULD resend the query using TCP on server port 53(decimal). This behavior is consistent with the current DNS specification [\[RFC1035\]](#).

- In future, EDNS0 is REQUIRED to send large packets.

Hence, the size of the UDP payload is no longer limited to 512 octets any more.

5.2 Standard sample

A resolver who wants to find the iDNs corresponding to an IP address

[1.2.3.4](#) would pursue a query of the form QTYPE=IPTR, QCLASS=IN, QNAME=4.3.2.1.IN-ADDR.ARPA, and would receive:


```

+-----+
Header   | OPCODE=SQUREY, RESPONSE, AA |
+-----+
Question | QNAME=4.3.2.1.IN-ADDR.ARPA.,QCLASS=IN,QTYPE=IPTR |
+-----+
Answer   | 4.3.2.1.IN-ADDR.ARPA. IPTR  "zh-cn" "name1-in-utf8" |
        | 4.3.2.1.IN-ADDR.ARPA. IPTR  "zh-tw" "name2-in-utf8" |
        | 4.3.2.1.IN-ADDR.ARPA. IPTR  "ja-jp" "name3-in-utf8" |
        | 4.3.2.1.IN-ADDR.ARPA. IPTR  "ko-kr" "name4-in-utf8" |
+-----+
Authority | ... |
+-----+
Additional | ... |
+-----+

```

6 Open Issues

1. API issues on the resolver side.
2. the granularity of the language info. (per domain name? per label? within label?)

Practically, we believe it is enough for the iPTR info to be expressed as |01|ELT|language|size|utf8|size|utf8|...|, meaning the LANGUAGE TAG is used to define the language of the Fully Quantified Domain Name. However, FQDNs could still exist in the form of "English-in-utf8.Chinese-in-utf8.English-in-utf8." And more than 1 language can exist in the same label. Should such level of detailedness be supported? Or a simple meta-type like "mixed-language" is enough?

3. If language info should somehow be relatable to an iDN itself(nothing to do with PTR...) and how?

As a suggestion, if a new RR TYPE INAME is established to relate iDN to current domain name, there will be two merit. One is we don't to do anything with PTR. Second is if we cache the INAME RRs to the DNS caches, then it can reduce the upper layer name servers' jobs. Actually, the feature of the new RR TYPE is quite similar to CNAME and DNAME, meaning name-to-name.

FQIDN: Fully Qualified Internationalized Domain Name.

Then the INAME RR is expressed following:

iDN	INAME	traditional domain name
-----	-------	-------------------------

About the first merit, When the client looks up an IP address to iDNs then the server will reponse not only corresponding PTR RR but corresponding INAME RRs to the client. Furthermore, the problem in LANGUAGE TAG can be avoided.

For example:

4.3.2.1.IN-ADDR.ARPA	PTR	traditional-domain-name
iDN-1	INAME	traditional-domain-name
iDN-2	INAME	traditional-domain-name

About the second merit, INAME is not only can be used in address-to-name mapping but name-to-address mapping.

For example:

traditional domain name	IN	A	host address
iDN-1		INAME	traditional-domain-name
iDN-2		INAME	traditional-domain-name

When the client looks up an iDN or traditional domain name to its corresponding IP address, if the server reponses not only A RR but INAME RRs to the client. And the client cache these RRs to its DNS cache. Then the next time, maybe some queries can be resolved in DNS cache.

References

[IDNE] Mrc Blanchet & Paul Hoffman, "Internationalized domain names using EDNS", [draft-ietf-idn-idne](#).

[RFC1034] P. Mockapetris, "DOMAIN NAMES - CONCEPTS AND FACILITIES", November 1987, [RFC1034](#)

[RFC1035] P. Mockapetris, "DOMAIN NAMES - IMPLEMENTATION AND SPECIFICATION", November 1987, [RFC1035](#)

[RFC1766] H. Alvestrand, "Tags for the Identification of Languages", March 1999, [RFC 1766](#)

[RFC1886] S. Thomson, C. Huitema, "DNS Extensions to support IP version 6", December 1995, [RFC1886](#)

[RFC2279] Francois Yergeau, "UTF-8, a transformation format of ISO 10646", January 1998, [RFC 2279](#).

[RFC2671] Paul Vixie, "Extension Mechanisms for DNS (EDNS0)", August 1999, [RFC 2671](#).

[ISO 639] ISO 639:1988 (E/F) - Code for the representation of names of languages - The International Organization for Standardization, 1st edition, 1988 17 pages Prepared by ISO/TC 37 - Terminology (principles and coordination).

[ISO 3166] ISO 3166:1988 (E/F) - Codes for the representation of names of countries - The International Organization for Standardization, 3rd edition, 1988-08-15.

Acknowledgements

James Seng has given many comments in our e-mail discussions.

Authors' Information

Hongbo Shi
Waseda University
3-4-1 Okubo, Shinjyuku-ku
Tokyo, 169-8555 Japan
shi@goto.info.waseda.ac.jp

Jiang Ming Liang
i-DNS.net
8 Temasek Boulevard
#24-02 Suntec Tower Three
Singapore 038988
jiang@i-DNS.net